Fakultät für Informatik

Potentials and Limitations of Contributive Social Capital Systems

Sebastian Schams

Vollständiger Abdruck der von der
Fakultät für Informatik
der Technischen Universität München zur Erlangung des akademischen Grades
eines Doktors der Naturwissenschaften (Dr. rer. nat.)
genehmigten Dissertation.

Vorsitzende: Prof. Dr. Anne Brüggemann-Klein

Prüfende/-r der Dissertation:

1. Priv.-Doz. Dr. Georg Groh

2. Prof. Dr. François Bry

3.

Die Dissertation wurde am 09.07.2018 bei der Technischen Universität München
eingereicht und durch die Fakultät für Informatik am 26.10.2018 angenommen.

POTENTIALS AND LIMITATIONS OF CONTRIBUTIVE SOCIAL CAPITAL SYSTEMS

SEBASTIAN SCHAMS

Doktor der Naturwissenschaften (Dr. rer. nat.)
Social Computing Research Group
Fakultät für Informatik
Technische Universität München

June 2018 – Version 1.0

LONG ABSTRACT

In recent years, we have been experiencing an increasing amount of online communication taking place on social media, social networking platforms, and other means of information exchange. While this ease of information sharing and staying in touch improves some aspects of life, it has also been accompanied by new challenges. One significant issue is the spread of lies and fake news on social media. Users also complain about the inability to correctly assess trustworthiness and expertise of other participants, which is especially relevant when confronted with new information.

Following the design science framework, this thesis introduces, discusses, and investigates a novel approach to mitigation of some of these issues by creating transparency regarding a person's contributive social capital (CSC). There is ongoing research about how individual pro-social characteristics can be analyzed and visualized. Additionally, there are implemented reputation and feedback systems on individual platforms and systems that analyze input from several data sources. The CSC system, which is conceptualized in this thesis, builds on this work, and aims to calculate a universal, possibility topic-dependent CSC value for each participant. CSC is a characteristic that describes a person's value-add to the overall social capital of a network and encompasses competence, trust, and social responsibility.

For the assessment of CSC, the system combines several evaluation methods. It is structured along three pillars which can be combined to create an individual score. The first pillar is the use of supervised learning algorithms for the analysis of social networking platforms like Facebook or Twitter and other data sources. The second pillar is built on a dedicated market system that creates and maintains CSC scores based on virtual currency transactions. The third pillar permits recognition of the CSC demonstrated in offline accomplishments with the help of endorsements and certifications.

Foundational elements of this system were investigated in several experiments. The investigations included the analysis of three major online platforms (Facebook, Twitter, and Quora), experiments on data from a scientometrics database, an agent-based simulation, and a dedicated experiment with 242 participants which was conducted over the course of nine weeks. The investigations were individually tailored to the respective data sources, e.g., by selecting appropriate features.

These experiments were compiled using a variety of additional analyses. Among others, subjectivity and biases in online communication and knowledge exchange platforms were examined, as well as potential use cases of the CSC system, and an implementation as a smart contract on the blockchain.

The investigations support some of the hypotheses upon which the CSC system has been conceived. The CSC scores extracted from the different data sources displayed a positive correlation with ground truth values. Additionally, the market system allowed the extraction of scores that reflect the ground truth to some extent. It is furthermore possible to visualize from which group individuals receives their support.

The experiments also demonstrated that the prediction of CSC is not yet possible with a high degree of certainty which limits the applicability. While we discuss how the system may be used to mitigate the issues stated in the beginning, a full-scale implementation and full-scale evaluation of such a system could not be part of these investigations due to resource restrictions.

Recent data scandals and the (mis)use of social network analysis methods for economic or political gains underline the importance of this research and give reason for concerns. To provide the full picture, we review some of these scandals and discuss the concerns and ways to mitigate the risk of misuse of the CSC system.

SHORT ABSTRACT

In recent years, we have been experiencing an increasing amount of online communication, which was accompanied by new challenges like the spread of fake news and the difficulty to assess someone's trustworthiness and competence.

Following the design science approach, we conceptualized a system for the mitigation of these issues by introducing transparency about a user's contributive social capital (CSC), a metric that encompasses competence, trust, and social responsibility. The CSC is assessed by a combination of inputs that include, among others, the supervised learning analysis of online data sources and a novel mechanism based on virtual currency transactions in a dedicated social capital market.

The foundations of CSC assessment were investigated in several experiments, e.g., with analyses of three online platforms (Facebook, Twitter, and Quora), studies about CSC extraction from a scientometrics database, an agent-based simulation, and an experiment with 242 participants. Additional investigations included subjectivity analyses in online communication and knowledge exchange platforms, as well as potential use cases of the CSC system, and an implementation as a smart contract on the blockchain.

These experiments support some of the hypotheses upon which the system has been conceived as they demonstrate that CSC estimation with the proposed methods is to some extent possible and that the provenance of CSC support can be visualized. We also discuss potential shortcomings and concerns and review ways to mitigate them.

KURZFASSUNG

In den letzten Jahren hat die Online-Kommunikation stark zugenommen, begleitet von einer Vielzahl neuer Herausforderungen. Trotz zahlreicher positiver Entwicklungen lassen sich dabei auch negative beobachten, wie die Verbreitung von "Fake News" oder die Schwierigkeit, die Kompetenz und Vertrauenswürdigkeit von Interaktionspartnern in sozialen Netzwerken richtig einzuschätzen.

Zur Adressierung dieser Probleme wurde im Rahmen dieser Doktorarbeit ein System konzipiert, mit dem Ziel, Transparenz hinsichtlich des beigetragenen Sozialkapitals (engl. Contributive Social Capital/CSC) einer Person zu ermöglichen. Unter CSC wird das Maß an Fachkompetenz, Vertrauenswürdigkeit und sozialem Verantwortungsbewusstsein einer Person verstanden. Das CSC-System zielt darauf ab, konkrete, individuelle und potentiell themen-spezifische CSC-Werte mit Hilfe einer Kombination mehrerer Bewertungsmethoden zu ermitteln. Zu diesen gehören unter anderem die Analyse verschiedener Online-Datenquellen mit Supervised Learning und ein neuer Mechanismus, der auf Transaktionen virtueller Währung in einem dezidierten Sozialkapitalmarkt beruht.

Dem Design Science Ansatz folgend, wurde das CSC-System nicht nur konzeptioniert, sondern auch seine Kernelemente, insbesondere bezüglich der Grundlagen der CSC-Erfassung, in mehreren Experimenten untersucht. Diese umfassten unter anderem Analysen dreier wichtiger Online-Plattformen (Facebook, Twitter und Quora), Studien zur CSC-Extrahierung aus einer Szientometrie-Datenbank, eine agentenbasierte Simulation und ein Experiment mit 242 Teilnehmern. Weitere Untersuchungen beinhalteten Subjektivitätsanalysen von Online-Kommunikations- und Wissensaustauschplattformen sowie mögliche Anwendungen des CSC-Systems und die Implementierung als Smart Contract auf der Blockchain.

Die Ergebnisse der Experimente stützen einige der Hypothesen, nach denen das CSC-System konzipiert wurde und bestätigen, dass die CSC-Ermittlung mit den vorgeschlagenen Methoden zu einem gewissen Grad möglich ist. Zudem lässt sich durch Visualisierung der CSC-Ströme nachvollziehen, von welchen Gruppen ein Individuum Unterstützung erhält. Abschließend werden eventuelle Unzulänglichkeiten und Bedenken bezüglich des CSC-Systems aufgezeigt, insbesondere hinsichtlich aktueller Datenmissbrauchsskandale, und Möglichkeiten diskutiert, diese zu mitigieren.

OWN PUBLICATIONS

Schams, S. and Groh, G. (2018). Social capital extraction from different types of online data sources. *submitted for publication*.

Schams, S., Hauffa, J., and Groh, G. (2018a). Analyzing a user's contributive social capital based on acitivities in online social networks and media. In *Companion of The Web Conference 2018*, pages 1457–1464. First International workshop on Online Social Networks and Media: Network Properties and Dynamics. ACM. https://dl.acm.org/citation.cfm?id=3191593.

Schams, S., Hauffa, J., and Groh, G. (2018b). Exploring market systems to visualize, build, and develop contributive social capital scores. *submitted for publication*.

ACKNOWLEDGEMENTS

When I started writing this thesis, I imagined writing the acknowledgements has to be a great feeling because it implies that all the hard work is successfully finished. Now that I have written over 200 pages about my research in the previous years, I feel kind of sad. Sad because it also means that a time is ending that was exciting, informative and above all beautiful — especially because of the following people:

- My supervisor PD Dr. Georg Groh, who deserves my greatest thanks. From the very beginning he challenged and encouraged me at the same time. I love thinking of the long discussions about how we can improve the world with a social capital system. His input immensely improved this thesis and he taught me a lot — about computer science and life.
- Prof. François Bry for being the second reviewer of this thesis and for valuable and detailed input during our discussions and interactions.
- Prof. Jörg Ott whose chair for connected mobility welcomed our social computing research group to the joint doctoral seminars and gave us many new insights.
- The other doctoral candidates of the social computing research group, especially Hanna Schäfer and Jan Hauffa, whose common interest in machine learning and social network analysis inspired many interesting debates.
- The 242 students who participated in our experiments and provided interesting insights in the social media surveys.
- The reviewers and workshop participants of The Web Conference 2018 in Lyon, who provided valuable information, new insights, and encouragement.
- All the students who joined my own little research group on social capital during their bachelor's theses, master's theses, and guided research projects: Simon Zettler, David Hauer, Monika Varshney, Christian Höfer, Valeriia Chernenko, Johannes Feil, Panagiota Revithi, Maximilian Schmidt, Rauf Zeynalov, Patrick Süß, Niclas Hirtle, and Daniel Clot. I really enjoyed being your supervisor and valued all of our discussions. It was an honor to guide you in the interesting time that is the completion of one's studies.

Thank you very much to all of you.

Finally, I would like to highlight the great teamwork of all the professors, postdocs, doctoral cadidates, and students who work together in the "Finger 5" of the informatics building in Garching. Not only the professional collaboration but also the friendly daily interactions, the "coffee breaks", and the joint events will be remembered fondly and missed greatly.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

LIST OF ABBREVIATIONS

ABMS – Agent-Based Modeling and Simulation
ACM – Association for Computing Machinery
ACT – Author-Conference-Topic (topic model for scientometrics)
API – Application Programming Interface
BFS – Breadth-First-Search
CART – Classification And Regression Tree
CBOE – Chicago Board Options Exchange
CME – Chicago Mercantile Exchange
CSC – Contributive Social Capital
CSCW – Contributive Social Capital Weight
DAO – Decentralized Autonomous Organization
DBLP – Digital Bibliography & Library Project (computer science library)
EU ETS – European Union Emissions Trading Scheme
GSD – Graph-based Sybil Detection
HITS – Hyperlink-Induced Topic Search (also known as Hubs and Authority)
JS – Jensen-Shannon Divergence
KL – Kullback-Leibler Divergence
LDA – Latent Dirichlet Allocation
nDCG – normalized Discounted Cumulative Gain
NLP – Natural Language Processing
OSN – Online Social Network
OSNEM – Online Social Networks and Media
P2P – Peer-To-Peer
PBoC – People's Bank of China
Q&A – Question and Answer (e.g., Q&A portal)
ReLU – Rectified Linear Unit
RDF – Resource Description Framework
RSS – Residual Sum of Squares
SCC – Social Capital Currency
SMOG – Simple Measure of Gobbledygook (measure of readability)
SNP – Social Networking Potential
SVR – Support Vector Regression
TUM – Technical University of Munich
XACML – eXtensible Access Control Markup Language

Part I

INTRODUCTION AND FOUNDATIONS

This part introduces the subject matter and presents basic concepts, algorithms, and background information that help to understand the thesis. The first chapter (1) motivates the area of research by reviewing current challenges in online communications. It also introduces our research questions, the methodology, and an outline of the thesis. The second chapter (2) provides an overview of the term social capital and presents definitions of other relevant terms. Different online data sources that account for the majority of online communication and knowledge exchange are introduced in the third chapter (3). Finally, different algorithms for graph analysis (4.1), machine learning (4.2), statistical analysis (4.3), and natural language processing (4.4) are reviewed in chapter 4 to provide a common notational and conceptual ground for the later chapters.

1

INTRODUCTION

1.1 INTRODUCTION AND MOTIVATION

In recent years, new technologies have created a multitude of new opportunities. Two ancient human dreams, access to nearly unlimited knowledge and the ease to stay in contact with people all around the planet have become reality with overwhelming speed. Due to notebooks, smartphones, and the Internet, people can access most of humanity's knowledge in seconds and communicate with each other at the click of a button, e.g., in online social networks and media (OSNEM). It is no surprise that this led to an increase of online communication (Lenhart et al., 2010), (Gabbriellini, 2014), (True, 2017).

However, not everything is positive about these new possibilities.

Some of the information people have access to are not trustworthy. There are substantiated, yet not fully established suspicions that fake news have influenced the outcomes of both the US presidential election of 2016 and of the Brexit referendum of the same year. In the months before the presidential election, the voters were, e.g., "[...] exposed to misinformation and fake news stories very often [...]" (Goodwin-Ortiz, 2017). The spreading of false information is sometimes due to malintent but sometimes simply because people do not correctly self-assess their own knowledge (Anson, 2018).

Even during disasters, around 90% of active Twitter users were likely to spread untrue messages, e.g., rumors during hurricane Sandy and the Boston Marathon bombings (Wang and Zhuang, 2018).

The analysis of Facebook activity provides a similar picture. According to a study, people in two of three groups on Facebook "indulged in deceptive behavior designed to self-promote or aggrandize the individual" (Underwood et al., 2011).

Social media usage has also been correlated with depression (Sidani et al., 2016), which may be due to technology-based social comparison and feedback-seeking (Nesi and Prinstein, 2015).

Another issue can be identified on the side of the platform operators, who do not necessarily treat all the data as confidential as it would be desirable. An example is the 2016 Facebook data scandal, where an external company harvested the data of up to 87 million Facebook users and used it for political advertisements (NY-Times, 2018).

To better understand how users of social media platforms experience online communication, we provided a questionnaire to several hundred IT students at TUM university, which was answered by 242 students. The students are experienced with social media; 94.2% use Facebook and 38.0% use Twitter.

76.9% of the participants agreed that it is difficult to identify fake news on social media platforms. When conversing with new or unknown people online, 74.0% found it difficult to assess whether or not they can trust them. Both of these results are in line with the previously cited research about mistrust and fake news on OSNEM.

Next, we asked the students about their previous experiences when trying to find information and knowledge online. While 55.4% of the students are usually able to satisfy their information need with a web search, 59.5% of the participants said they found it difficult to identify an expert who could directly help them. This is an important issue for many students, as 63.6% answered they often have topics they would like to discuss with an expert.

When trying to help others, the students also faced problems. They were generally willing to provide help to friends or other people with questions in their field of expertise (91.3%). However, 80.6% found it difficult to advertise their expertise on social networking platforms.

Our goal is to contribute to the research that aims at improving this situation in general and communication on social media and social networks in particular. Openness and transparency are often seen as tools to decrease corrupt behavior (Bertot et al., 2010). Therefore, we investigate ways to create transparency in a way that is beneficial for all users, incentivizes supportive and social behavior, and allows people to advertise their skills.

There are already several approaches to create some level of transparency. Many platforms implement feedback and reputation mechanisms. Buyer and seller platforms like Amazon or eBay allow users to directly rate each other (Bhattacharjee and Goel, 2005). Facebook, Twitter, and other online communities implement ways to provide feedback about specific contributions, e.g., in the form of a "like", a retweet, or up- and down-voting. Additionally, there is some research about how trust, reputation, expertise, and other properties can be extracted from social networks (e.g., (Ziegler, 2009), (Yang et al., 2010), and (Su et al., 2012)), even with input from several platforms (Rao et al., 2015).

However, to the best of our knowledge, there is no holistic approach that tackles all of the discussed issues and their relations.

We conceptualized a system that creates transparency for users over several platforms and includes additional methods to allow an assessment that is as comprehensive as possible. Our research approach was structured in three steps.

The first step was the definition of a personal characteristic that, when visualized, allows a quick assessment of the person's value-add to the network. To address most issues identified with online communication, the characteristic should reflect several aspects of a user's persona:

- Knowledge and expertise ("the goal of communication is the transfer of information[...]" (Hauser, 1996). Therefore, people with high knowledge and expertise can provide information that may be of interest to other network participants).
- Trustworthiness (without trust, one cannot believe the shared information).
- Social responsibility (only a user who behaves in a social way provides help and information to the other network participants).

A term that encompasses these aspects and that has been used to describe social networks in prior work is *social capital*. We refine this term in section 2.1 and introduce the term *contributive social capital* (CSC) as an explicit metric in the context of this work.

The second step was to conceptualize a system that puts a value to this metric. This system includes three main pillars that represent different assessment methods to allow a holistic evaluation. The first pillar is the analysis of interactions on different social networking and social media platforms, as well as academic citation networks. The second pillar is a market system that builds topic-specific CSC scores based on transactions of virtual currency that are related to contributions to the network. And thirdly, endorsements and certifications allow to include CSC that has been demonstrated in offline life into the system. Motivations for these pillars, as well as further details and specific design choices are discussed in part II of this thesis.

An experimental investigation of the foundations of CSC assessment, upon which the system is built on, was the third and most comprehensive step of our research approach. We conducted a variety of experiments to examine the extraction of CSC from various online data sources and to evaluate to what extent market systems can be used for an additional assessment. In one of these experiments, 242 participants interacted on a social networking platform, participated in a virtual market, and provided a ground truth approximation in the form of mutual assessments. It was not the goal of this thesis to fully implement the CSC system but to provide fundamental investigations as part of a design science circle.

The experiments about CSC assessment are complemented by a variety of additional investigations, e.g., regarding the assessment of biases and subjectivity in online communication, potential risks and shortcomings of the system, as well as a potential implementation on the blockchain.

This thesis is guided by a variety of research questions that are listed in the next section 1.2. The research methodology is described in section 1.3. Several bachelor's and master's theses and articles were created in the context of this thesis, as described in section 1.4. Finally, section 1.5 provides an outline of the thesis.

1.2 OVERARCHING RESEARCH QUESTIONS

With the help of these experimental studies, as well as literature research and analyses, the following research questions were investigated. These questions define the direction of our research on a high level. They are further detailed and extended in the chapters that describe the corresponding experiments.

1. What is a suitable personal characteristic that can be tracked and visualized to improve the experience of online communication?

2. How could a system look like that automatically measures the contributive social capital of a person in a holistic way and maintains these scores over time?

3. How can an individual CSC assessment for each user be extracted from previous interactions on OSNEM and other data sources?

4. How can a market system be used to build CSC scores of the market participants based on their transactions?

5. How is it possible to contextualize the CSC of a person along different topics?

6. How do participants experience the market system?
7. What are use cases of the CSC system for individuals, companies, and governmental institutions?
8. How can a CSC system be implemented in a secure and transparent way?
9. What are potentials and limitations of the CSC system?
10. Can the system be expanded to visualize subjectivity and biases of individuals regarding different topics?

The first research question directly starts where section 1.1 left off at the definition of a metric that can be used to create transparency in online communication. As discussed, (contributive) social capital is a suitable personal characteristic. The definition of this term and the motivation for its use are provided in chapter 2. The second research question asks for the theoretical concept of a system that can be used for the assessment of this characteristic. The concept of this system and the experimental investigation of foundational elements of CSC assessment are important contributions of this work. The extraction of CSC estimates from user activity and feedback on OSNEM and other data sources is the focus of research question 3. Online social networks and media account for an important part of online communication. Additionally, we investigated CSC assessments on Q&A portals and academic citation networks. The possibility of a market-based CSC assessment is introduced by research question 4. This is investigated with experiments in which CSC scores are modeled based on currency transactions in a dedicated market. The remaining research questions 5 to 10 further refine the system and its implementation.

The findings with regard to the research questions are summarized in chapter 20.

1.3 RESEARCH METHODOLOGY

This thesis uses two research methodologies: empirical research and design science.

Empirical research describes the process of deducing knowledge from empirical observations. These observations can then be used to support or disprove an hypothesis. We used empirical research, e.g., in the context of social network analysis.

Design science is an approach that provides specific guidelines for research about information systems. The seven guidelines introduced by Hevner et al. (Hevner et al., 2004) are listed below and discussed with regard to this thesis.

- **Design as an Artifact.** Hevner et al. state that "design science research must produce a viable artifact in the form of a construct, a model, a method, or an instantiation" (Hevner et al., 2004). In the course of this thesis several such artifacts were created to investigate different aspects of the CSC system. One example is the social capital market system, which is investigated with an experiment (chapter 14) and a market simulation (chapter 15).
- **Problem Relevance.** Design science's goal is to provide "technology-based solutions to important and relevant business problems" (Hevner et al., 2004). In section 1.1, we discussed a variety of problems in online communication. The

CSC system aims to promote transparency and trust, and to create additional motivations for altruistic and pro-social behavior, which may address some of these issues.

- **Design Evaluation.** "The utility, quality, and efficacy of a design artifact must be rigorously demonstrated via well-executed evaluation methods" (Hevner et al., 2004). The foundations of the CSC system were evaluated with several experiments, especially regarding the first two CSC assessment pillars. The experiments were analyzed thoroughly, e.g., with statistical measures.

- **Research Contributions.** We discuss various manners to estimate how the interactions of a person in various social networks are perceived by that person's peers. These assessments are set in context to contributive social capital and embedded into the CSC system. The system and especially the estimation methods were investigated with several experiments. A nine-week user study with 242 participants indicated that CSC can be measured in two ways: via extraction from social networking platforms as well as by using market systems for a continuous CSC assessment. These investigations were complemented by large scale-studies on Facebook, Twitter, Quora, and a scientometrics database, as well as a market simulation. All contributions are summarized in chapter 20 with regard to the overarching research questions.

- **Research Rigor.** Following Hevner et al., "design-science research relies upon the application of rigorous methods in both the construction and evaluation of the design artifact." (Hevner et al., 2004) Previous related work was reviewed in detail (see, e.g., chapter 10) to ensure that the system implements the most relevant techniques (e.g., supervised learning with neural networks or random forests). All results were evaluated with common statistical tools (chapter 4.3).

- **Design as a Search Process.** The search process can be interpreted as an iterative circle in which new insights are used to improve the results of the next iteration. The investigations of this thesis can be seen as (part of) a first circle of such a process and lay the basis for a full implementation of the system. The conclusions with regard to the potentials and limitations of CSC systems are summarized in chapter 21.

- **Communication of Research.** Some aspects of this work were already published, as listed on page vii. This thesis can be regarded as documentation of all results.

Based on these guidelines, design science can be represented as a circle, as depicted in figure 1. After motivating the problem (step 0), an iterative process begins that includes the design of a concept or a system that has the goal to solve or tackle the problems (step 1). With the help of experiments (step 2) this system is tested and after evaluation (step 3) the concept is refined.

For parts of this thesis behavioral science is also important. It takes the needs of people and organizations into account and thereby complements the technical information systems research. The Facebook data scandal (see chapter 17.3) has illustrated

Figure 1: Visualization of the design science approach

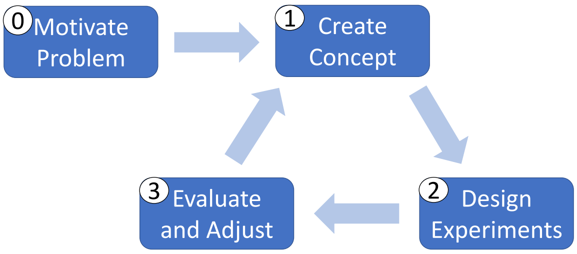

how influential social network analysis is and will continue to be. Therefore, it is essential to not only regard the technological side, but the human implications with all relevant advantages and disadvantages. Figure 2 provides a conceptual framework (Hevner et al., 2004) that combines behavioral-science and design-science paradigms.

Figure 2: Information systems research framework (Hevner et al., 2004)

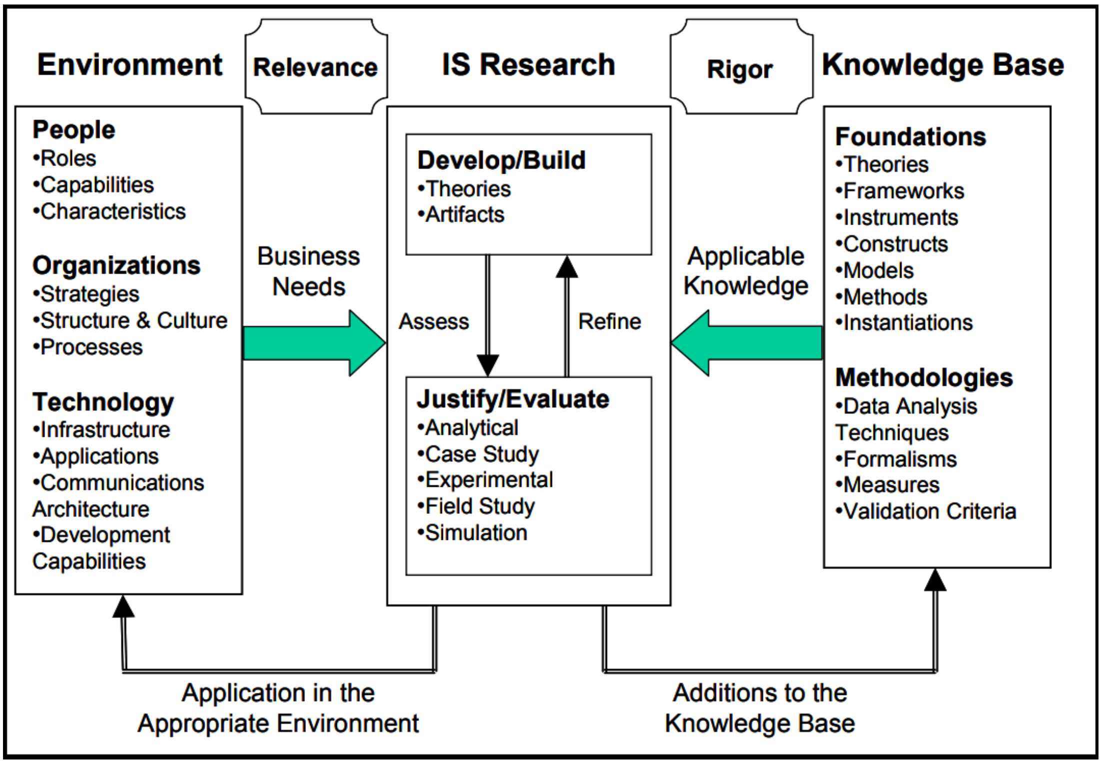

In the center of figure 2 is the information systems (IS) research approach. In our case this is the design, evaluation, and refinement of a contributive social capital system. Our investigations are based on theoretical foundations and computer science methodologies that are introduced in part I of this thesis. This is represented on the right side of the figure. The applications of the CSC system with regard to people, companies, and governmental institutions – which is illustrated on the left side of figure 2 – is discussed throughout this thesis and separately in chapters 16 and 17.

1.4 SUPERVISED THESES AND PUBLICATIONS

Several bachelor's and master's theses and a guided research project were supervised in the context of this thesis. A complete list of all theses is provided in appendix D.

A list of all publications can be found on page vii. Sebastian Schams was the lead author of these articles. Some parts of this thesis are direct excerpts from these papers. To improve the reading flow of this thesis, the quotations are not denoted individually. Instead, appendix E provides a detailed and complete list of all excerpts.

1.5 OUTLINE OF THE THESIS

This thesis is structured in six parts.

Part I introduces the topic and reviews the theoretical background upon which this thesis is built. It encompasses the definitions of social capital and contributive social capital, related terms, and an overview of different online data sources that we investigated in the experiments. The technical foundations cover graph analysis, supervised machine learning algorithms, statistical analysis, and natural language processing.

Based on these foundations, part II introduces the concept of a CSC system. Its goal and basic functionality are described in chapter 5. It consists of three pillars that are described in chapters 6, 7, and 8. The combination of the three pillars, the system architecture, and design choices are discussed in chapter 9. In the following two parts, the first two pillars are investigated experimentally and several experiments that explore their implementation are discussed.

Part III focuses on the extraction of CSC from online data sources. It reviews previous work (chapter 10) and discusses five experiments. The social capital experiment (chapter 11) was a nine-week study with 242 participants who participated in an online social networking (OSN) platform and provided ground truth data. The other four experiments were conducted on excerpts from Facebook, Twitter, Quora, a scientometrics data set and are described in chapter 12.

Part IV describes our investigations about the use of market systems to assess CSC scores. After a review of existing market mechanisms (chapter 13), two experiments are described. The experiment presented in chapter 14 represents the market side of the social capital experiment. All 242 participants had access to a market platform and could exchange currency. Based on these monetary transfers, CSC scores were created. These investigations are complemented by a market simulation with 1,000 agents, which is described in chapter 15.

The experiments described in the previous parts indicate that contributive social capital can be measured — at least with regard to the limitations of the experiments.

Part V describes a variety of other investigations that complement the previous experiments. At first, several different use cases for the CSC system are discussed (chapter 16). Especially with regard to some data misuse scandals in the recent years, it is important to consider potential risks of the system. This discussion takes place in chapter 17, at the hands of two examples. An interesting addition to the CSC system is the visualization of subjectivity. In some topics there are two or more distinct sides (e.g., Democrats and Republicans in US politics). Subjectivity investigations aim to visualize the provenance of support of a group or individual. This may allow to identify potential biases or people with well-balanced views. We investigated subjectivity with regard to contributive social capital in five different experiments: on Twitter, Facebook, a scientometrics data set, Quora, and the social capital experiment data. These investigations are summarized in chapter 18. A way to securely implement a social capital market on a large scale in a transparent and trustworthy way is to use blockchain technologies. This possibility is investigated in chapter 19.

In the last part of this thesis (part VI), the results are summarized following the overarching research questions, conclusions about the potentials and limitations of CSC systems are drawn, and areas for future work are discussed. This is rounded off with the appendix in part VII that provides further information about the experiments and the supervised theses.

2

SOCIAL CAPITAL

2.1 SOCIAL CAPITAL AS A METRIC

One goal of our research is to investigate ways to mitigate issues in online communication, as presented in section 1.1. There is a lot that is unknown about new online conversation partners, including their competence, trustworthiness, and motivations. Consequently, it is a valid assumption that improving transparency may solve these issues (Bertot et al., 2010). If the transparency also sheds a light upon the social behavior of network participants it may additionally motivate altruistic behavior.

In face-to-face encounters, people are used to making assessments of their counterparts in fractions of a second (Wargo, 2006). In order to approximate this experience in an online setting, one needs to create clarity with a single *metric* that is easy to understand. Such a metric – e.g., in the form of a score – could be quickly interpreted and provide a first impression. This score could additionally be split up along topics to provide further insights once the initial assessment is completed. This contextualization by categories is discussed in section 9.2.8.

There is a variety of personal characteristics that may act as such a metric. Some, like expertise, trustworthiness, or reputation have been researched extensively in the context of social media (see, e.g., (Golbeck, 2008) and chapter 10 of this thesis) and certainly address some of the issues. However, they focus primarily on single aspects and do not provide a holistic view of the individual. A characteristic that implements aspects of many of these other terms and is overall focused on the benefit of the whole network is social capital. In this chapter, we discuss different definitions, introduce a refined definition for the context of this work and set social capital in context to other characteristics.

2.1.1 *Social Capital*

There is a variety of definitions for social capital. Robison et al. attribute this to the highly context-dependent nature of social capital and argue that social capital has often been defined with a specific application in mind (Robison et al., 2002). In general, one can differentiate between two types of social capital. The first describes the properties of social networks on a macro level. An exemplary definition in this context was offered by Putnam, who describes social capital as "features of social organization such as networks, norms, and social trust that facilitate coordination and cooperation for mutual benefit." (Putnam, 1995) Alternatively, one can look at social capital from the perspective of an individual (ego level) and describe their surrounding micro network. Lin, for example, describes social capital as "a broad concept usually focused on the values obtained by being part of a social network and thus, referred to as the sum of social resources." (Lin, 2002) This is in line with the earlier definition by Nahapiet and Ghoshal who describe social capital as "the sum of the actual and potential resources embedded within, available through, and derived from the network of relationships possessed by an individual or social unit. Social capital thus comprises both the network and the assets that may be mobilized through that network." (Nahapiet and Ghoshal, 2000)

The choice to select social capital as a metric for our system is backed up further by (Fukuyama, 2001), who concludes that "social capital is important for the efficient functioning of modern economies and stable liberal democracy."

2.1.2 *Contributive Social Capital (CSC)*

For the purpose of our research we are interested in a metric on the individual level, opposed to the macro level described in the previous section. Visualizing the social capital a user has access to via their relationships and network would be a way to measure social capital on the individual level. However, it does not achieve the intended purpose of creating transparency regarding the contributions and helpfulness of individuals. To achieve this goal, we need to examine the opposite perspective: the social capital a user adds to their network. To better describe this metric, we introduce the term *contributive social capital* (CSC).

In order to track and visualize contributive social capital as a metric in the system, we need to define it and assess what other characteristics CSC comprises of. A better understanding of these aspects facilitates a ground truth measurement for the experiments, e.g., by asking for peer assessments of the respective CSC constituents.

To achieve the goals of the system, CSC should be defined in a way that combines two objectives: reflect the value-add of a person to their network in online communication and target the issues identified in section 1.1. Many personal characteristics can be associated with (online) communication. The goal of communication is the transfer of information (Hauser, 1996). Consequently, all character traits that encompass knowledge, expertise, experience, or reflect other ways of information access are of relevance. To convey information a variety of soft skills and additional characteristics are of importance. These encompass, but are not limited to, the person's trustworthiness, reputation, benevolence towards others, attractiveness, eloquence, or even physical strength. While it is the goal to create a holistic representation of each participant, it is not feasible to merge all of these characteristics into the single metric CSC, as it would complicate the ground truth assessment. Additionally, the communication of CSC scores and future research would be facilitated by a definition that is focused on the essentials. Characteristics of little importance should consequently be disregarded. Based on these considerations and the two initially stated objectives, we can define the three most important aspects of a person's CSC in online communication:

- Factual competence,
- trust and reputation, and
- social responsibility.

The dimension of factual competence comprises the knowledge, experience, and information a person brings to their network. In many online interactions this is the most important part of value-add, simply because the transfer of information is the main goal of communication (Hauser, 1996). It covers, for instance, reviews of products on sales platforms, answers on Q&A portals, programming help in the respective forums, or sharing of personal news on social networking platforms.

Being able to trust the information provider is an additional condition to regard the expertise of a person as value-add for the network — the expertise of an untrustworthy person does not necessarily increase the network's overall social capital. This also addresses, to some extent, the goal of the system to stop the spread of lies and fake news. Trustworthy people are less likely to spread such false information and can be more easily identified due to the CSC assessment provided by the system.

Finally, the dimension of social responsibility takes the willingness of an individual to help others and provide support and information into account. Without this active participation in the social network, there is no practical value-add for the whole network. This also reflects the goal of the system to promote pro-social and helpful behavior.

These three aspects reflect the majority of a user's value-add to the overall social capital of the network. Creating transparency regarding the CSC of network participants – especially if it is divided by topics – could, therefore, address the issues described in 1.1 and motivate altruistic behavior because people are ranked highly if they are competent, trustworthy, and willing to help others.

2.2 RELATED TERMS

As a measure for an individual's value-add to the social network, contributive social capital is related to several other characteristics. In this section, reputation, trust, expertise, and influence are defined in the context of OSNEM and set in relation to CSC.

2.2.1 *Reputation*

Hoffman et al. define *reputation* as "the opinion of the public towards a person, a group of people, an organization, or a resource" (Hoffman et al., 2009). This is consistent with other definitions (e.g., compare Jøsang et al. (Sang et al., 2007) who quote the Concise Oxford dictionary: "Reputation is what is generally said or believed about a person's or thing's character or standing"). In the context of social networks this means that the reputation of a participant is built with information that is available to the social network. Reputation can thus be defined as the sum of all subjective judgments of actors in a network regarding a single individual. From these definitions it becomes apparent that reputation is a subjective element, a degree of belief to which others intersubjectively believe in another person's competence in a certain field.

2.2.2 *Trust*

An even more individually subjective construct is *trust*. People can choose to trust or distrust others without considering their respective public reputation. Diego Gambetta's definition from 1988 is often used throughout literature: "trust (or, symmetrically, distrust) is a particular level of the subjective probability with which an agent assesses that another agent or group of agents will perform a particular action" (Gambetta, 1988). Jøsang et al. (Sang et al., 2007) contrast this definition with a quote from Mcknight et al. on what they call "Trusting Intention": "the extent to which one party is willing to depend on the other party in a given situation with a feeling of relative security, even though negative consequences are possible" (Mcknight and Chervany, 1996).

Similar to reputation, trust is often topic-focused and situation-dependent. The average person would, for example, trust a math teacher more with answering a mathematical problem than a person chosen at random. This aspect of trust is relevant for online interactions, especially when users primarily care about a single topic. A seller on an online auctioning site is, for instance, mainly interested in the buyer's willingness to pay. On the other hand, the reader of a threaded discussion board wants to know how competent the author of certain posts is in the topic and whether he or she can believe the recommendations.

2.2.3 *Expertise*

If we move from the example of buyer-seller networks to content-oriented internet platforms (e.g., threaded discussion boards or expert exchanges), the *expertise* of an information provider becomes increasingly important. According to Bozzon et al., expertise is closely related to someone's knowledge in a defined domain. They describe identifying experts within a social network as the task of "ranking the members of a social group according to the level of knowledge that they have about a given topic" (Bozzon et al., 2013). After the ranking, the top k experts may be chosen and selected to answer questions or fulfill tasks. Therefore, their expertise can be regarded as the ability to answer questions on certain topics with the help of their knowledge or experience. Expertise is in general verifiable based on facts, whereas reputation and trust are more subjective and less specific judgements.

2.2.4 *Social Influence*

A widely studied research subject is *social influence*. According to J. Sun, "social influence refers to the behavioral change of individuals affected by others in a network" (Sun, 2011). With the emergence of online social networks that document a wide range of user interactions, it became possible to study influence regarding a wide range of aspects and fields. Anger et al. conclude that only a small percentage of all users on Twitter are what they call "influencers." (Anger and Kittl, 2011) These people publish content that is in turn read and reposted by many followers. Identifying these people is of a high interest to companies for marketing and directed advertisements (Sun, 2011) and has been investigated, e.g., on Twitter (Hauffa et al., 2016).

2.2.5 *Overlap Between CSC and Related Terms*

CSC is defined as a person's value-add to their network. This value is primarily influenced by that person's relationships to others and by the information and resource access that person has. Therefore, all attributes of a person that facilitate social interactions, foster relationships, and promote information access can be regarded as aspects of (contributive) social capital.

Reputation can increase a person's CSC, as other people are more likely to connect and interact with, as well as provide information to a reputable person. This was investigated by Resnick et al. (Resnick et al., 2000) on online buyer-seller networks. On eBay, they observed that "sellers with stellar reputations may enjoy an extra premium on their services — a premium that users may be willing to pay for the security and the comfort of high quality services." If customers are willing to pay a premium price for a perceived better reputation, it is reasonable to assume that a higher reputation also correlates to non-monetary benefits. Additionally, the reputation is a measure for the work already invested into a network by the respective participant. In the case of researchers, e.g., reputation can be measured with the h-index (Hirsch, 2005), which directly indicates the amount of contributions by an individual that were regarded as important by their peers.

A similar reasoning can be applied to *trust*. Valenzuela et al. argue based on Putnam's research, "social trust facilitates associative behavior, fosters a strong civil society, and makes political institutions and officials more responsive [...]" (Valenzuela et al., 2009). This can be understood as an increase of CSC.

The *expertise* of a person directly correlates with the information and resource access aspect of CSC (see, e.g., (Recuero et al., 2011)). It increases the amount of information potentially made available to the network by a person. People who are close to this individual can now access this information, e.g., via discussions. This leads to a similar effect we have seen for reputation: The overall increase of social capital in the network can be attributed to the individual CSC of the person.

There is also a significant overlap between a user's *influence* and his or her CSC. Influence usually stems from several factors. Among those are the average novelty of a user's contributions (Keller and Berry, 2003), their eloquence (Keller and Berry, 2003), and the quality of the content they produce (Carmel et al., 2012). These factors act as local multipliers on the CSC of a person and increase its basic constituting components directly.

2.3 WAYS TO INFER CONTRIBUTIVE SOCIAL CAPITAL

To best describe a system for the assessment of CSC, we first need to understand how people assess each other and how modern technology can be leveraged to potentially improve this process.

There are several ways to infer characteristics of a person in general and their contributive social capital in particular. The first way is the most straightforward. In every face-to-face interaction we judge our counterpart based on verbal and non-verbal communication. These assessments take place in less than a second (Wargo, 2006). Nonverbal, implicit behavior is especially important "in situations in which unfamiliar persons interact, with one seeking to influence the other (as in political speeches or advertising)." (Mehrabian, 2017) This direct assessment usually takes place automatically in everyday interactions and does not allow to gain an impression of the other person before the interaction occurs.

The second way to infer CSC is based on interactions a person has had in the past. In the age of social media, a variety of information about social media users is available. This data consists of activity on social media, like posts on Facebook, tweets on Twitter, or contributions to other networking platforms. People sometimes even disclose more information about themselves on Facebook than in everyday interactions (Christofides et al., 2009). Based on this past behavior one can form an opinion about other users and their CSC.

A third way to assess characteristics of a person are direct feedback systems. These are often already implemented on social media platforms and can be analyzed in a similar fashion. Examples for such feedback mechanisms are the "likes" or favorites on Facebook and Twitter, star-based review systems for buyers and sellers on Amazon, or up-vote and down-vote mechanisms as employed in communities like Quora and Reddit. These systems incorporate the feedback of others and help us to build an opinion about the usefulness of postings and indirectly the CSC of contributors.

Another way to obtain information about character traits of others, is to ask for assessments of people who dealt with them in the past. This is a similar idea to the web-of-trust approach (Stark, 2001). It can be combined with a type of PageRank (see section 4.1.2.5) or Hubs-and-Authority (see section 4.1.2.6) approach that values a person's feedback based on their own social capital.

Finally, there is a mechanism that our economy is based on: a market system with free trade and its own currency. Market systems were rarely used for the assessment of character traits — or solely for those directly correlated with wealth. However, it might be possible to implement a new, virtual currency that is linked to specific personal traits. Interactions with this currency could then be used to infer the respective feature. This hypothesis is supported by the success of prediction markets (Wolfers and Zitzewitz, 2004). We investigate this concept in the context of social capital.

The system for the assessment of the CSC of individuals is based on the observations described in this section. It is explained in detail in part II of this thesis.

3

ONLINE DATA SOURCES AND CONTRIBUTIVE SOCIAL CAPITAL

The first pillar of the CSC system is the assessment of a person's CSC based on data from different online data sources. There is a variety of platforms that allow people to communicate and interact online. In this section, five popular data sources are briefly described at the hands of examples. The focus is put on sources that allow people to communicate, share information, or help each other. Other platforms (e.g., pure gaming communities) that are of less relevance for social capital, are disregarded.

When discussing the functionality of the different data sources, we also describe what feedback mechanisms are in place and how they are related to contributive social capital.

3.1 MICROBLOGGING

Twitter is an example of a microblogging (Java et al., 2007) service with over 300 million active users per month and the self-declared mission "to give everyone the power to create and share ideas and information instantly, without barriers"[1]. Twitter messages have a maximum length of 280 characters and can contain text, URLs, pictures, mentions (references to other users with the symbol @ and their name), hashtags (which is a type of label that is usually connected to the post's context or a current event), and locations. Mentions can increase the popularity of the mentioned user, however, this popularity is not explicitly stated. Hashtags can increase the visibility of tweets because interested users can search for certain hashtags. Participants can follow others to receive their tweets directly (Weng et al., 2010), which is like a subscription. The connection between followers is often based on reciprocity, which was demonstrated by Weng et al., who found that "80.5% of users have 80% of users they are following follow them back" (Weng et al., 2010). If a user agrees with a post it is possible to "retweet" the full comment or to quote the text and add a personal statement. Both acts increase the original post's visibility as it is made accessible to a larger crowd (namely the followers of the retweeter). It has been demonstrated by Hofer and Aubert that there is a correlation between Twitter usage and social capital because Twitter allows members to stay in touch with many people (Hofer and Aubert, 2013). In their study, participants who spent more time on Twitter, also reported a higher perceived online bridging social capital. Additionally, Hofer and Aubert investigated how the numbers of followers and followees correlate to a user's online social capital and found a positive correlation as well: People with more followees/followers usually have a higher bridging/bonding social capital (Hofer and Aubert, 2013). As the time spent on Twitter correlates with a user's social capital, it is reasonable to assume that the other metrics that signal a user's engagement correlate with their (contributive) social capital as well. One of these indicators could be the

1 https://about.Twitter.com/company (retrieved 2017-06-30)

number of posts a user published. Indicators that include the perspective of other users might be even more useful: for example, the number of retweets users generate or the times they are mentioned in other posts. Users can, however, be mentioned by others who disagree with them fully. This is an interaction that increases the number of times they are mentioned, but it does not necessarily correlate with a high CSC.

Analysis methods that implement and go beyond these single performance indicators are discussed in section 10.1. An experimental investigation of CSC extraction from Twitter is described in section 12.6.

3.2 ONLINE SOCIAL NETWORKING PLATFORMS

Social networking platforms are online services that allow their members to connect with each other. Facebook is the most popular social networking platform and has over 2 billion active users[2]. The numbers of friends, "likes", received and posted messages can be used as indicators for a user's importance for the network and can in turn be regarded as connected to this user's CSC. Burke et al. argue that "it is plausible that creating and consuming undirected messages, allowing users to keep in touch, will lead to increases in social capital" (Burke et al., 2011). Similar to the previous example of Twitter, the passive indicators (received "likes", friends, comments, and responses to own posts) are more trustworthy indicators for a person's value within the network because they incorporate feedback from others and are, therefore, more complex to manipulate.

Previous work about assessing user characteristics based on interaction on online social networking platforms is discussed in section 10.2. The extraction of CSC from Facebook was investigated in section 12.5.

3.3 DIRECT COMMUNICATION

There are many ways to directly communicate with another person online. Microblogging and social networks often implement forms of direct communication by allowing users to send messages to each other. Even more straightforward direct conversation is the use of mobile messengers or email. A popular mobile messenger is WhatsApp[3]. For the purpose of CSC identification email has certain advantages. The length of a message is restricted neither by the service provider, nor by the ease of input, which is less comfortable for long texts on a smartphone. Also, it can include personal as well as professional information and, therefore, a more complete portrayal of someone's CSC is possible. As this data is usually private, there are no built-in indicators to visualize social capital to the public. However, there are data sets available that allow an investigation of social capital and its components in this context (e.g., Enron email data[4]). Metrics that can be used for analysis are the number of outgoing and incoming emails as well as the message content (see section 10.3).

2 https://de.newsroom.fb.com/news/2017/06/facebook-bedankt-sich-bei-2-milliarden-monatlich-aktiven-menschen/ (retrieved 2018-04-19)
3 https://www.whatsapp.com/ (retrieved 2018-04-19)
4 Cohen, W. Enron email data set, http://www- 2.cs.cmu.edu/ enron/(retrieved 2018-04-19)

Previous work about the extraction of CSC related character traits from direct communication is discussed in section 10.3.

3.4 scientometrics

Scientific citation networks and co-citation networks are formed by publications that reference other papers. Kas et al. (Kas et al., 2012) list several of their properties. Citation networks are directed, acyclic networks where the bulk of the network is static and only the leading edge is dynamic (Kas et al., 2012). Jøsang et al. also mention another important property: there are only positive referrals (Sang et al., 2007). Therefore, it is not easily possible to sanction authors whose publications were highly cited in the beginning but later turned out to be controversial.

Compared to the previous networks, citation and co-citation networks evolve slowly and there is no real exchange of directed communication content between people. In the context of social capital extraction, however, the fact-based scientific nature offers distinct advantages. The content is reviewed before it is published, which guarantees a high level of quality. Therefore, performance metrics can be used to draw conclusions about the expertise of authors and co-authors.

The direct performance metric in citation networks is the number of citations an author or paper receives by other published articles.

Section 10.4 reviews previous work about the analysis of authors in citation networks and section 12.7 describes an experiment we conducted to assess contributive social capital in citation networks.

3.5 threaded discussion boards and q&a portals

Threaded discussion boards organize the discussions of their users by topics in threads. The website Reddit is a popular representative of such a community and one of the most frequented websites in the world[5]. Reddit is a social news website with discussion threads (Weninger et al., 2013) and is organized along different "Subreddits" that focus on specific topics (e.g., news or the NBA). Users can either link other online content and briefly summarize it with a headline, or post thoughts of their own. To promote the best content, Reddit has a voting mechanism that allows registered users to up- or down-vote content. The resulting score (related to the difference between up- and down-votes a user receives) is called *karma* and comes in two flavors, link karma and comment karma (Bergstrom, 2011). The first one is awarded for links to other websites (e.g., news articles) that the community regards as worth reading, and the latter is assigned to a user's comments for either its informative content, its originality, or other aspects, like humor.

Another news and threaded discussion board is Slashdot (Lampe et al., 2007). The content posted on Slashdot is often related to science, technology, or politics and summarized by the posters. The posts are organized along different topics, and users can comment on them to start discussions. The voting system is slightly different from Reddit. Only randomly assigned moderators can up- or down-vote posts, not every registered user.

5 http://www.alexa.com/siteinfo/Reddit.com (retrieved 2018-04-19) displays a global rank of 6.

Question and Answer (Q&A) portals are quite similar in their setup: People can register and then converse with each other. As the name suggests, the focus is set on posing questions and other users answering them. A popular Q&A portal is Quora[6]. Stack Overflow[7] is another well-known Q&A portal that also displays characteristics of threaded discussion boards. Its main audience are software engineers and it contains a variety of computer related knowledge. Similar to Quora and Reddit, readers can up- and down-vote questions and answers.

Previous work about the analysis of CSC related user characteristics on these platforms is discussed in section 10.5. Section 12.8 describes the experimental assessment of CSC on Quora.

6 https://de.quora.com/ (retrieved 2018-04-19)
7 https://stackoverflow.com/ (retrieved 2018-06-14)

4 TECHNICAL FOUNDATIONS

In this chapter, theoretical foundations are presented to provide a common notational and conceptual ground for the later chapters. The discussed methods and algorithms were used for the analyses and their interpretation. Graph analysis (section 4.1) was used to extract information about the users of social networks, e.g., about their importance for the overall network, or to identify different communities within the OSNs. The supervised machine learning algorithms described in section 4.2 were used to investigate the extraction of CSC from different data sources. Section 4.3 briefly explains statistical analysis methods used to interpret results. Natural language processing, which is discussed in section 4.4, was used as input for clustering and to identify features for the machine learning analyses.

4.1 GRAPH ANALYSIS

Social networks, including those created by the communities described in section 3, can be represented as graphs. For this purpose, the network participants (e.g., users on Facebook) are nodes which are connected by edges that describe the nature of their relationship, e.g., friendships or sent messages. Network analysis methods are used to identify nodes that connect different node clusters, have a high influx of edges, or are in other ways important for the graph. This information can then be transferred to the people corresponding to these nodes.

There are several methods to analyze graphs to obtain these insights. The ones important for this thesis can be roughly categorized in two groups: clustering and centrality measures.

4.1.1 *Clustering*

An important application of unsupervised machine learning is the partitioning of data items or feature vectors into different groups, called clusters (Jain et al., 1999). If an item is assigned to exactly one cluster the method is called hard clustering. In some cases it may be useful to assign an item to several clusters, e.g., when clustering scientists by the fields in which they published. This is called soft or fuzzy clustering. Over the course of this thesis, we used several different clustering methods: metric clustering, graph clustering, attribute-based clustering, and topic-based clustering.

4.1.1.1 *Metric Clustering*

Metric clustering is the most popular type of clustering. All items that are to be clustered are represented with feature vectors that contain all features along which the grouping should happen. Then, the algorithm clusters the results by either finding an appropriate distance or similarity measure between the feature vectors. A well-known metric clustering algorithm is *k-means*. It divides the items into a predefined number k of clusters by minimizing the quadratic distance of the feature vectors to the centers of the respective clusters. Each object is then iteratively assigned to the closest cluster centroid. K-means usually converges quite fast. A detailed review of metric clustering can be found in (Jain et al., 1999).

4.1.1.2 *Graph Clustering*

As the name suggests, graph clustering is aimed at identifying related groups of vertices in graphs (Schaeffer, 2007). It can be achieved with different algorithms. Divisive algorithms start with one single cluster that contains all nodes and divide this cluster into subclusters. Agglomerative approaches do the opposite: They start with one cluster for every single node and successively merge them to larger ones. The output of both methods is a hierarchy of clusters, which is also called dendrogram and visualized in figure 3. (Schaeffer, 2007)

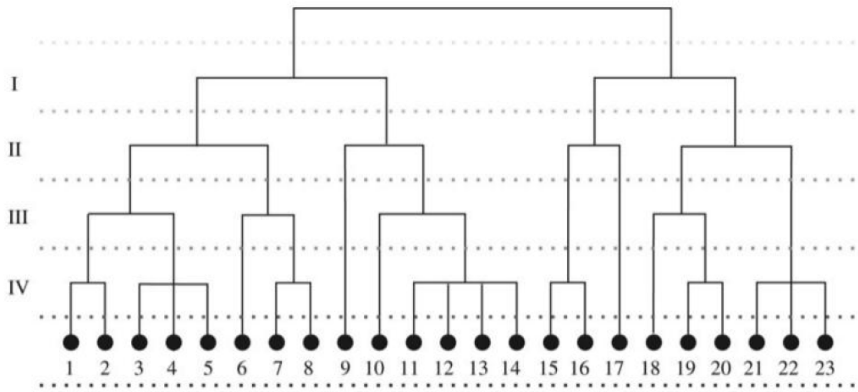

Figure 3: Visualization of hierarchical clustering (Schaeffer, 2007)

4.1.1.3 *Attribute-Based Clustering*

The most straightforward way to identify clusters in a network is to select an attribute and group all items according to the values of this attribute. This has the intuitive advantage that the clusters can be easily understood and communicated. Additionally, it allows clustering of discrete-valued attributes, as the following examples illustrate. The users of Facebook or Twitter can, e.g., be clustered by gender (female and male) and researchers in citation networks by university or conference. For real-valued attributes one needs to define values for the clustering. On Facebook or Twitter those thresholds could correspond to a user's activity level (less than 1 contribution a day, 2-5 contributions, more than 5 contributions), or age.

4.1.1.4 *Topic-Based Clustering*

Topic-based clustering describes all clustering based on topical information. Let us assume a topic distribution $\theta \in [0,1]^K$, $\sum_{i=0}^{K-1} \theta^{(i)} = 1$ that describes for each user in which topics the user is interested in. The goal of the clustering is to identify groups of users with similar topical interests or activities. Based on the three previously presented methods we can define different approaches for topic-based clustering.

- **Metric clustering:** Metric clustering methods can be applied by using the topic distributions of each item (e.g., users or documents) as feature vectors. As a distance measure one can use the Jensen-Shannon divergence (see section 4.3.4).
- **Graph clustering:** For a social network one can create a graph with the users as nodes and the similarity between their topic distributions as weight of the edges between them. This similarity can be measured with the Jensen-Shannon divergence. By applying graph clustering to the resulting graph, one can create topic-specific clusters. This was successfully applied by Aker et al. to cluster reader comments on online news platforms (Aker et al., 2016).
- **Attribute-based clustering:** One can attribute each person to one of the K topics by assigning each user to the one topic which the user is most interested in. Alternatively, one could define a threshold and attribute each user to all topics above this limit. This would be soft clustering as it can assign several topics to a single user.

4.1.2 *Centrality Measures*

With the help of centrality measures the most important nodes or edges in a graph can be identified. There is a variety of centrality measures that usually identify slightly different nodes as important, which is due to their different focus. In the following, degree, betweenness, and closeness, as well as eigenvector, PageRank, and HITS centrality are briefly discussed. Additional centrality measures and variations of the ones described in this section can be found in (Koschützki et al., 2005).

4.1.2.1 *Degree Centrality*

The well-known degree centrality is defined as

$$c_i^{DEG} = deg(i). \tag{1}$$

c_i^{DEG} is the number of nodes directly connected to node i. If the edges are, e.g., the friendships on Facebook, the c_i^{DEG} of person i is equal to their number of friends. In undirected graphs, e.g., the follower/followee relationships on Twitter, one can additionally distinguish between in-degree and out-degree of a user. (Sun, 2011)

4.1.2.2 *Closeness Centrality*

Another popular measure to assess the connectedness of a person in a network is to calculate the average shortest path to all other nodes in the network:

$$c_i^{CLO} = e_i^T S \mathbf{1}, \tag{2}$$

where e_i is a column vector whose ith element is equal to 1 and all other elements are 0. The element S_{ij} of matrix S is the shortest path from node i to j and $\mathbf{1}$ is a vector whose elements are all 1. The smaller c_i^{CLO}, the shorter all average distances to the other nodes are. (Sun, 2011)

4.1.2.3 *Betweenness Centrality*

The betweenness centrality describes how important a node is for connecting other nodes within the network. Freeman's betweenness centrality is defined as

$$c_i^{BET} = \sum_{j,k} \frac{b_{jik}}{b_{jk}}, \tag{3}$$

where b_{jk} is the count of shortest paths between nodes j and k and b_{jik} is the number of times this shortest path runs through the investigated node i. (Sun, 2011)

4.1.2.4 *Eigenvector Centralities*

The idea of eigenvector centralities is that the importance of a node does not only come from the respective connections of the nodes but from the importance of the connected nodes as well. This can be visualized with the following equation

$$c_i^{EIG} = a \cdot \sum_j g_{ij} c_j^{EIG}, \tag{4}$$

where a is a constant and g_{ij} is 0 if node i and node j are not connected and 1 if they are. To calculate c_i^{EIG} we, therefore, have to calculate c_j^{EIG}. To solve this, one can write the equation in matrix form:

$$\mathbf{C} = a\mathbf{GC}, \tag{5}$$

which for $\frac{1}{a} = \lambda$ becomes the eigenvector equation

$$\mathbf{GC} = \lambda\mathbf{C}, \tag{6}$$

where $\mathbf{G}$ is the auxiliary matrix and λ is the eigenvalue. There are, in general, many different eigenvalues λ for which an eigenvector solution exists. With the help of the additional requirement that all entries of the eigenvector are non-negative, a centrality measure can be determined. This is due to the Perron-Frobenius theorem which states that there is one unique largest real eigenvalue for any irreducible matrix. A sufficient condition for this is that the graph is strongly connected. This largest eigenvalue results in the centrality measure. (Aggarwal, 2011)

A high eigenvector centrality score means that the node is connected to many other nodes that themselves have high scores.

4.1.2.5 *PageRank Centrality*

The PageRank algorithm (Page et al., 1999) was originally invented by Larry Page and Sergey Brin and used as basis for prioritizing websites in the search engine Google. Wu et al. (Wu et al., 2007) summarize the functionality of PageRank:

> "In essence, PageRank relies on the democratic nature of the Web by using its vast link structure as an indicator of an individual page's quality. It interprets a hyperlink from page x to page y as a vote, by page x, for page y. Additionally, PageRank looks at more than just the sheer number of votes or links that a page receives. It also analyzes the page that casts the vote. Votes cast by pages that are themselves important weigh more heavily and help to make other pages more important."

This method can also be used for other types of graphs, e.g., with the users of social networks instead of web pages as nodes. The idea behind PageRank is similar to that of eigenvector centrality, as it also regards the importance of the incoming node.

4.1.2.6 *Hubs and Authority*

The Hubs and Authority algorithm, which is also known as Hyperlink-Induced Topic Search (HITS), was introduced by Jon Kleinberg (Kleinberg, 1999). Following (Schütze et al., 2008), the HITS algorithm works as follows. For every query, each node in the network is assigned two scores: authority score and hub score. Originally these nodes were websites, in the context of CSC analysis the nodes become network participants. People who are sources of information on the topic of the query are described as authorities, people who have many connections to authorities, but are themselves no authorities are called hubs. This leads to a circular definition that amounts to an eigenvector definition: prime authorities are linked to by prime hubs, and prime hubs are link to prime authorities. These scores can then be computed in an iterative fashion.

4.2 SUPERVISED LEARNING ALGORITHMS

There are many applications for which it is difficult to devise algorithms that solve the problem very well, even though many data points are available for the analysis (Alpaydin, 2016). An example for these applications are translation tasks. Supervised machine learning algorithms that "learn" directly from the data can be used in these cases and are often able to deliver better results than human programmers (Alpaydin, 2016). As input, these algorithms require a data set that is labeled with ground truth values. For a prediction task, the ground truth is, e.g., the correctly predicted value for a set of features from the input data set.

In this thesis we investigate the use of supervised machine learning algorithms for the prediction of contributive social capital scores from participants of online social networking platforms. The algorithms and evaluation mechanisms used, are briefly described in this section.

4.2.1 *Linear Regression*

Simple linear regression aims to describe the relationship of the dependent variable Y and the independent variable X in the form

$$\hat{y}_i = w_1 x_i + w_0 \tag{7}$$

The intercept w_0 and the coefficient w_1 were estimated based on previously known or measured pairs of X and Y and are the same for all x_i and $\hat{y}_i$. Linear regression can also be expanded by adding a quadratic $+w_2 x_i^2$, or higher dimensional term. (Alpaydin, 2010)

One can also use multiple explanatory variables with different coefficients w_j. This is called *multiple linear regression* and can be represented in matrix notation as:

$$\hat{\mathbf{y}} = \mathbf{X}\mathbf{w}, \tag{8}$$

where $\hat{\mathbf{y}}$ is the vector of predicted values, $\mathbf{X}$ is the matrix of row vectors $\mathbf{x}_i$, and $\mathbf{w}$ is the parameter vector. When we mention linear regression during this thesis, we usually refer to multiple linear regression.

The difference between the predicted $\hat{y}_i$ and the correct value y_i is the prediction error. The model coefficients w_j are estimated by minimizing the sum of the squared prediction errors $\sum_{i=1}^{n}(y_i - \hat{y}_i)^2$. There are some assumptions that need to be met for the linear regression to be valid (following (Gelman and Hill, 2006)):

- The residuals are independent from one another.
- The variance of the errors is equal (also called homoscedastic).

There is no consensus in literature on whether the assumption that the errors should be normally distributed is necessary (Gelman and Hill do not recommend it (Gelman and Hill, 2006)). Section 4.2.7 presents methods to evaluate the goodness of the fit, which also reflect these assumptions.

4.2.2 *Decision Tree Regression*

Decision trees can be used for regression and as classification models. Both tasks are achieved by a series of if-then-else decision rules that are learned from training data. Each of the decision rules breaks down the data set into subsets and thereby incrementally builds a tree. The resulting tree has decision nodes and leaf nodes. A decision node has two or more branches that reflect the value of the attributed test (e.g., if the decision is "height $\leqslant$ 1.7 meters" then there are two branches, one for people taller than 1.7 meters and one for all others). Leaf nodes represent the end of a path and the value associated with it is the output predicted by the tree. (Alpaydin, 2016)

Several algorithms can be used to construct a tree, among them ID3 (Quinlan, 1986), its extension C4.5 (Quinlan, 1993), and CART (Breiman, 2017). These algorithms work recursively and minimize the overall error of the regression. This is achieved by comparing a metric before and after each potential split, i.e. before implementing a new decision rule. This metric is usually the Gini index (Gini, 1912), the information gain (entropy), or the residual sum of squares RSS, which is the error term of the left and right branch after the split and is defined by the following equation

$$\mathrm{RSS} = \sum_{\mathrm{left}} (y_i - \bar{y}_L^*)^2 + \sum_{\mathrm{right}} (y_i - \bar{y}_R^*)^2, \tag{9}$$

where $\bar{y}_L^*$ is the mean value of y in the left node and $\bar{y}_R^*$ the mean value of y in the right node. Once the error after a split is acceptable, i.e. below a pre-defined threshold, or the maximum depth of the tree is reached, a leaf node is created. The mean value of the output of all instances arriving at that node is calculated and assigned as prediction value to the leaf. If the error lies above the threshold, the data reaching the respective node is split further — in a way that the sum of the errors in the branches is minimized. This procedure is continued until every branch ends in a leaf. In this way a piecewise constant approximation of the prediction with discontinuities at leaf boundaries is created. (Alpaydin, 2010)

Once all the rules of the decision tree are determined, $\hat{y}$ can be directly predicted for new data. A distinct advantage of decision trees is that its decision rules can be easily understood and communicated. However, decision trees are prone to over-fitting and often have a high variance, in a way that small changes in the data can result in completely different trees. Also, greedy algorithms cannot guarantee to return the globally optimal decision tree. To prevent over-fitting, one can use *pruning* to "cut back" unnecessary splits that mainly model the noise of the data. Pruning is achieved with cross-validation (see section 4.2.6).

4.2.3 *Random Forest Regression*

Tin Kam Ho first introduced random forests, an ensemble technique that works by combining multiple decision trees (Ho, 1995).

The first step of the algorithm is to grow individual trees on independent samples of the training data, by drawing N cases at random with replacement. Following the procedure presented in section 4.2.2 a decision tree is now trained on the data. This process is repeated until the number of trees is reached, which can be specified by the user and determined experimentally (Raschka, 2015). The default number of trees in scikit (Pedregosa et al., 2011) learn is 10.

The predicted value $\hat{y}$ is the average of the predicted values of all separate trees (for regression) or the majority vote (for classification).

Due to this combination of different decision trees, the risk of over-fitting is reduced and the results are usually better than those of decision trees. The transparency of the decision making, however, is lost.

4.2.4 *Neural Networks*

Innovations in all disciplines of human engineering have been inspired by nature. One such idea, which was reinvented in the late 1980s, is using artificial neural networks for data analysis (Alpaydin, 2010). Originally inspired by cognitive processes happening in the human brain, artificial neural networks are able to produce astonishing results in various applications ranging from stock market predictions to disease prevention.

A neural network consists of neurons, the basic processing elements. The input x_i of the neuron comes either from the environment or from other neurons and comes with a weight w_{ij} that is determined during the training of the neural network. In the simplest case the prediction of an output neuron y_j is the weighted sum of the inputs:

$$y_j = \sum w_{ij} x_j + w_0, \tag{10}$$

with w_0 being the intercept value that makes the model more general.

Figure 4: Illustration of a feed-forward neural network (da Silva et al., 2016)

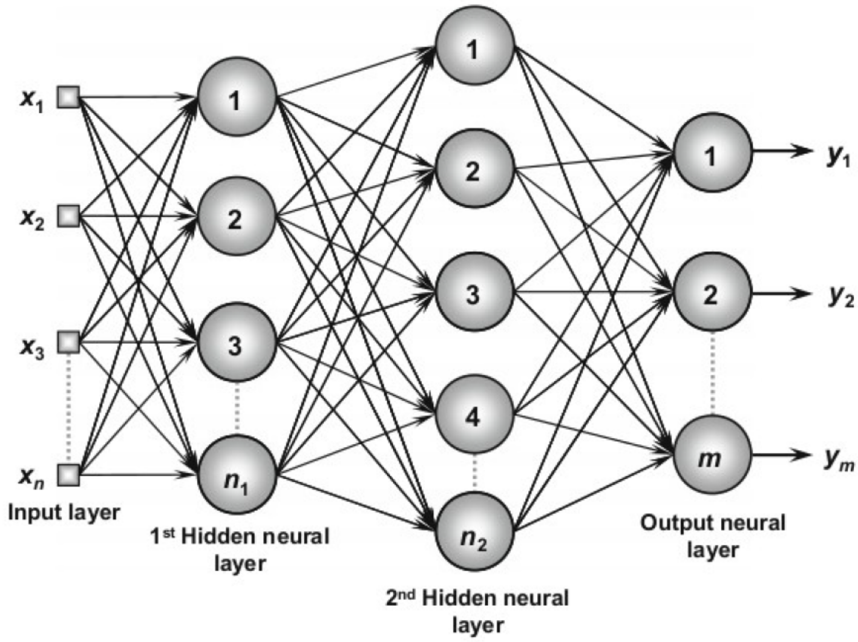

By combining neurons in layers, one can create a larger neural network, as illustrated in figure 4. On the left side we find the input layer, where all the features are fed into the neurons of the network. There are several hidden layers (in this case two) and an output layer where the predictions y corresponding to the features are returned. Building on equation 10 we can now describe the input I_j of a neuron j in layer L as dot product of the outputs $y_i^{(L-1)}$ and the weights $w_{ij}^{(L)}$ of the edges connecting neuron i in layer $L-1$ with neuron j in layer L:

$$I_j^{(L)} = \sum_{i=0}^{n_L]} w_{ij}^{(L)} y_i^{(L-1)} \tag{11}$$

The output value of the neuron is then computed by applying an activation function g:

$$y_j^{(L)} = g(I_j^{(L)}) \tag{12}$$

Typical activation functions are the sigmoid (logistic function), the hyperbolic tangent function, or nowadays rectified linear functions (ReLU). The weights w_{ij} of all edges are determined via back-propagation of the error of the prediction through the network (da Silva et al., 2016).

Once a neural network is trained in this way, it is pretty fast and usually provides predictions with a high quality. A disadvantage is, however, that the reasoning behind a specific neural network prediction is a "black box" that makes a complete analytical understanding difficult, e.g., how small parameter adjustments influence the output.

4.2.5 *Support Vector Machines*

Support vector machines (SVM) are another type of supervised learning algorithm that can be used for both classification and regression.

Every object from the training data, which is defined by several features, is represented by a feature vector. For classification of the objects it is the task of the support vector machine to define a hyperplane that separates all the vectors of both classes. Additionally, the distance of the vectors that are closest to the hyperplane and the hyperplane is maximized. The plane that maximizes this margin is called *optimally separating hyperplane* and allows for a better classification of new data. To define this hyperplane, only the support vectors that are closest to the plane are required. If no linearly separable hyperplane can be defined to separate the different classes, one can introduce slack variables or map the vector space to a higher dimensional space in which a linear separation of the different sets is possible. The latter can be achieved with suitable kernel functions that represent the hyperplane in higher dimensions but do not massively increase computation time during the transformations. (Alpaydin, 2016)

Support vector regression (SVR) works in a similar fashion. The goal of SVR is not to define a hyperplane that separates the different classes, but to construct a hyperplane that passes near each vector in a way that all vectors lie within a specified distance of the hyperplane. For the construction of the hyperplane only the subset of the vectors on the margin are required and the vectors that lie close to the model prediction are disregarded. In the two-dimensional space the resulting hyperplane is a line which is, in most use cases, non-linear.

The accuracy of SVMs made them a popular choice for a variety of applications, e.g., in bioinformatics and language processing. They share the "black box" problem of neural networks and random forests, when it comes to analytically understanding the results (Alpaydin, 2016).

4.2.6 *Cross Validation*

A problem that may occur especially when using higher order supervised machine learning algorithms is *over-fitting*. That means that the algorithm fits the training data (almost) perfectly, but is no longer flexible in regard to new data input that may follow slightly different patterns.

To investigate whether the algorithm is over-fitting, the available data is often split into two sets a training set and a test set. The algorithm is then trained on the training set and evaluated on the test set to make sure that the results generalize to an independent data set.

K-fold cross-validation is a variation of this method that is especially useful when there is only a small data set available and one does not want to exclude some of the data for mere testing purposes. The principle of 10-fold cross-validation is explained in the following. At first, the data is split into ten parts of equal size. Then, the algorithm is trained with nine parts as training data and evaluated on the 10th part. For the evaluation one can, e.g., use the mean average error (see section 4.2.7). This procedure is repeated nine times until every part was used once as test set and nine times in the training data. The average of the evaluation metric (in this example the mean average error) of all ten runs is then calculated. In this way the whole data set can be used to train the algorithm.

4.2.7 *Evaluating the Quality of the Fit*

There are several methods to evaluate the quality of the fit and the performance of regression models. In the following, we have a closer look at the coefficient of determination R^2, errors, residual plots, and QQ-plots.

The coefficient of determination R^2 is a statistical measure that indicates how much of the variance in Y is predictable from the independent variable. It is defined as

$$R^2 = \frac{\sum_{i=1}^{n}(\hat{y}_i - \bar{y})^2}{\sum_{i=1}^{n}(y_i - \bar{y})^2}. \tag{13}$$

$\bar{y} = \bar{\hat{y}}$ is the mean of all y_i, $\hat{y}_i$ the predicted values and y_i the true values. R^2 lies between 0 and 1 and can be interpreted as measure of how well the estimated values $\hat{y}_i$ approximate the real data points.

The error of the fit, i.e. the difference between the correct value y and the value predicted by the algorithm $\hat{y}$, can be used to compare the quality of the prediction or the performance of different algorithms. Typically, one uses the mean average error of all predicted values or the root mean squared error, which puts more weight on larger discrepancies.

Residual plots are created by plotting the predicted values $\hat{y}_i$ on the x-axis and the residuals on the y-axis. Based on the resulting distribution one can draw several conclusions about the regression model. If the residuals are evenly distributed on both sides of $y = 0$, there are likely no biases and the average expected value of all residuals is presumably 0. All other distributions in the residual plots, e.g., curvatures, funnels, or periodic patterns are a sign of a deviation of the classical regression model assumptions. A curvature (e.g., all residuals lie on one side of $y = 0$ and their value grows with increasing $\hat{y}_i$) is a sign for missing higher order terms in the regression model. A funnel shape (e.g., residuals are small for low values of $\hat{y}_i$ and increase rapidly for higher values of $\hat{y}_i$) indicates heterogeneous variances, which contradicts the model assumption of equal variance of errors. (Yan, 2009)

Finally, quantile-quantile-plots (also known as QQ-plots) are a measure to compare two probability distributions. This can be used to investigate whether the residuals originate from a normal distribution. The QQ-plot is created by plotting the ordered observed and standardized residuals on the y-axis and the ordered theoretical residuals (in this case from a normal distribution) on the x-axis. The values of the quantiles of both distributions are sorted from smallest to largest. At last the ordered data is combined to pairs and plotted. If the residuals follow a normal distribution the resulting scatter plot should be a straight line. Deviations from the straight line indicate that the residuals do not follow a normal distribution. (Heiberger and Holland, 2004)

4.3 STATISTICAL ANALYSIS

A variety of statistical tools and methods were used for the evaluation of the experiments. In this section the most important concepts are briefly presented. A more in-depth discussion can be found in the cited literature.

4.3.1 *Pearson Correlation*

The Pearson correlation coefficient (also called Pearson's r) is a dimensionless measure to describe the degree and direction of the linear correlation between at least two variables. For two sample data sets $\{x_1, ..., x_n\}$ and $\{y_1, ..., y_n\}$ it is defined as

$$r = \frac{\sum_{i=1}^{n}(x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum_{i=1}^{n}(x_i - \bar{x})^2}\sqrt{\sum_{i=1}^{n}(y_i - \bar{y})^2}}, \tag{14}$$

with $\bar{x} = \frac{1}{n}\sum_i x_i$ and $\bar{y} = \frac{1}{n}\sum_i y_i$.

The resulting value r lies between -1 and +1. -1 describes total negative correlation, +1 total positive correlation, and 0 no linear correlation, which means that $\{x_1, ..., x_n\}$ and $\{y_1, ..., y_n\}$ are (almost) completely independent (Runkler, 2015). Pearson's r requires x and y to be normally distributed. It can only be used to determine the correlation of linear relationships.

4.3.2 *Spearman Correlation*

While Pearson's r is limited to describing linear relationships, the Spearman rank correlation coefficient (also called Spearman's ρ) can describe all monotonic relationships between two variables x and y. (Awad and Khanna, 2015)

Spearman's ρ calculates the Pearson correlation of the ranked values of the variables. In the first step $\{x_1, ..., x_n\}$ and $\{y_1, ..., y_n\}$ are converted to ranks, $\{r_{x,1}, ..., r_{x,n}\}$ and $\{r_{y,1}, ..., r_{y,n}\}$, and then ρ is defined as:

$$\rho = 1 - \frac{6\sum_i(r_{x,i} - r_{y,i})^2}{n(n^2 - 1)}. \tag{15}$$

The interpretation of Spearman's ρ is similar to Pearson's r. -1 describes total negative and +1 total positive correlation, which means that the variables are perfect monotone functions of one another. A ρ of 0 indicates that there is no correlation, i.e. the relative variable y does not increase or decrease relative to x.

4.3.3 *Statistical Significance Analysis*

With the help of the different correlation coefficients, which were described in the previous sections, one can investigate whether two paired sets of data are related. For the interpretation of the statistical relevance of the correlation, not only the values of Pearson's r or Spearman's ρ are relevant, but also the number of pairs in the data. The reasoning is that a correlation coefficient close to +1 which was calculated on a small sample is not as significant as the same value on a larger sample.

The standard method for the assessment of the significance of empirical analysis is the *p-value*. It represents the probability that the correlation coefficient would have arisen in case the null hypothesis is true. For correlation analysis, the null hypothesis is usually that x and y are unrelated. This can be clarified with an example. Let us assume that the correlation coefficient between two data sets of 100 values was calculated to be 0.254. This leads to the p-value of 0.01 and means that there is only a 1% chance that one can make these observations if the variables were unrelated. If a correlation analysis on an even larger data set results in the same correlation coefficient, the p-value would be even smaller. (Fenton and Neil, 2012)

4.3.4 *Jensen-Shannon Divergence*

Similarity measures are used to compare two probability distributions p and q, e.g., to investigate how similar two topic distributions of two different data sets are. One popular measure for this comparison is the Jensen-Shannon (JS) divergence, which is based on the Kullback-Leibler (KL) divergence. Following (Lee, 1999), the Jensen-Shannon divergence is defined as follows:

$$JS(p,q) = \frac{1}{2}[KL(p||m) + KL(q||m)], \text{ with } m = \frac{p+q}{2} \tag{16}$$

$KL(p||m)$ is the Kullback-Leibler divergence, which is defined as:

$$KL(p||q) = \sum_{i=1}^{n} p_i \log \frac{p_i}{q_i}. \tag{17}$$

Both of these measures calculate the distance between two distributions. To infer the similarity from these measures one can use $1 - JS(p,q)$.

4.3.5 *Test for Normality*

To test whether a variable is distributed following a normal distribution, Shapiro and Wilk introduced the Shapiro-Wilk-Normality test (Shapiro and Wilk, 1965). The null hypothesis H_0 assumes that the data was drawn from a normal distribution. The test statistic W is defined as

$$W = \frac{(\sum_{i=1}^{n} a_i x_{(i)})^2}{\sum_{i=1}^{n}(x_i - \bar{x})^2}, \tag{18}$$

where $x_{(i)}$ is the ith smallest number in the sample, $\bar{x}$ is the mean of the sample, and the constants a_i are defined as $(a_1, ..., a_n) = \frac{m^T V^{-1}}{(m^T V^{-1} V^{-1} m)^{1/2}}$, with $m = (m_1, ..., m_n)^T$ being the expected values of the order statistics of a normal distribution and V the covariance matrix of the expected order statistics. (Shapiro and Wilk, 1965)

4.4 NATURAL LANGUAGE PROCESSING (NLP)

4.4.1 *Language Analysis and Simple Features in NLP*

We analyzed the interactions of people in OSNEM as well as the publications of scientists in academic networks. The goal of this analysis was to extract features and predict individual social capital scores with the help of supervised learning algorithms. Some of the features used for this analysis were created with language analysis.

Straightforward features are the length of the posts, comments, or articles written, which is simply the count of characters. Another feature is the number of times a certain letter or symbol is used, e.g., the number of "#" on Twitter corresponds to the number of hashtags used by a person.

Another language feature that we used is the SMOG index. SMOG stand for Simple Measure of Gobbledygook and was introduced by (Mc Laughlin, 1969). Based on the complexity of the words used by a person it can estimate the number of years of education required to understand the text. The $SMOG(T)$ level of a text T is defined as follows:

$$SMOG(T) = 1.043\sqrt{\frac{30 \cdot \phi}{\sigma}} + 3.1291, \tag{19}$$

where ϕ is the number of words with three or more syllables and σ is the number of sentences in the text. The standard error of $SMOG(T)$ is 1.5.

Two comprehensive surveys about language analysis are provided in (Van Hout and Vermeer, 2007) and (Lu, 2012).

4.4.2 *Topic Modeling*

Identifying which topics a person, e.g., an author in an academic citation network, is most involved with offers interesting insights for social capital analysis and allows clustering of groups along their main interests.

Probabilistic topic models can be used to achieve this task. A popular state-of-the-art model is Latent Dirichlet Allocation (LDA) which was introduced in 2003 (Blei et al., 2003).

Following (Blei et al., 2010) and (Steyvers and Griffiths, 2007), we briefly present the methodology of LDA.

All available text can be seen as a set of documents. Each of these documents can be described with a distribution over topics, which are mentioned within the document. Each topic is modeled as a distribution over words. The topic distribution of document d can then be described as $\theta_d \in [0,1]^K$, where K denotes the number of topics and $\sum_{k=1}^{K} \theta_d^{(k)} = 1$. The word distribution of each topic is defined in a similar fashion. If W is the number of all words in the vocabulary, the word distribution of topic k is $\phi_k \in [0,1]^W$ with $\sum_{w=1}^{W} \phi_k^{(w)} = 1$.

With a topic distribution θ_d and word distributions ϕ_k for each topic $k \in 1, ..., K$, a document d of length N can be created with the following generative process:

For each word w_n of the N words:

1. choose a topic $z_n \sim \text{Multinomial}(\theta_d)$
2. choose a word $w_n \sim \text{Multinomial}(\phi_{z_n})$

The left side of figure 5 illustrates this process. Topic 1 deals with money and financial institutions and topic 2 with waterways. Words occur multiple times indicating their importance within the topic. Some words, like bank, can appear in both topics.

There are three documents that were created with words from the two topics. The topic distributions θ_d of the documents are represented by the arrows. Document 2 has, e.g., a topic distribution θ_{d2} of (0.5,0.5). It was created by selecting one of the topics with probability 0.5 and then choosing a word from the topic with the respective probability.

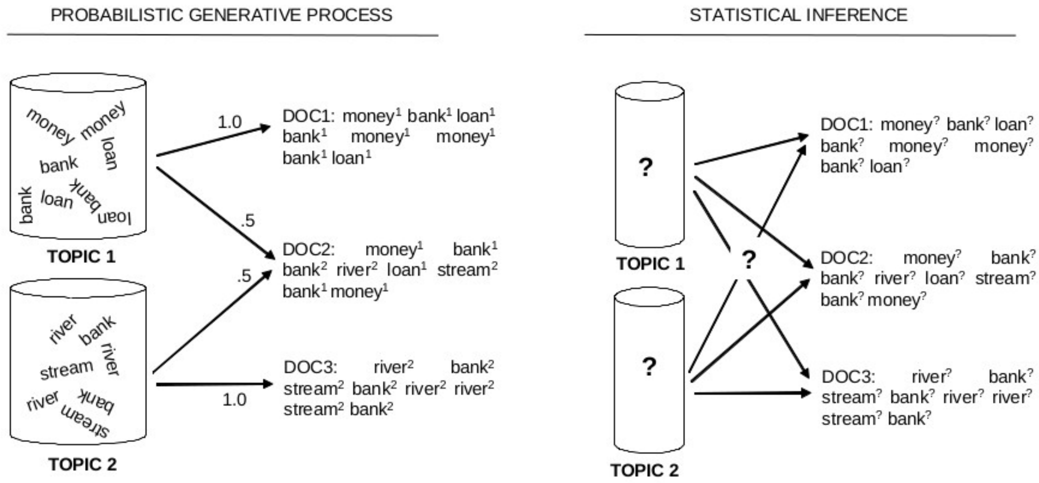

Figure 5: Illustration of probabilistic topic modeling at the hands of an example (Steyvers and Griffiths, 2007)

In general, the topic distribution θ_d and word distributions ϕ_k are not known in advance. This is also the case for our applications. Instead of creating documents from existing distributions, we want to infer the topic distributions from a set of documents. I.e. the algorithm should extract the latent variables (the topic distributions θ_d of all documents d and the word distributions ϕ_k for all topics k) in a way that they best describe the documents. This process is illustrated on the right side of figure 5.

The inference algorithm used in this thesis is based on (Hoffman et al., 2010).

4.4.3 *String Matching*

String matching algorithms identify matching text segments within longer text strings. In the course of this thesis, these algorithms were used to identify the university a scientist belongs to. The university name was often not stated explicitly but within a larger string. There are two general types of string matching: precise and approximate matching. We used approximate string matching to allow small deviations which may be caused by spelling mistakes. The algorithm that we used for the analysis is based on the Levenshtein distance (Levenshtein, 1966). It is a distance measure for strings that is defined by the number of operations needed to transform one string into another (e.g., inserting, substituting, or deleting a character).

A comprehensive review about approximate string matching is provided in (Navarro, 2001).

Part II

CONCEPTUALIZATION OF A CONTRIBUTIVE SOCIAL CAPITAL SYSTEM

This part of the thesis presents the concept of a contributive social capital system that allows to assess a person's contributive social capital weight based on input from different data sources. Section 2.3 discussed basic observations about the determination of characteristics of people, especially their social capital. Based on these observations, a system was envisioned to assess, build, and maintain indicator scores for contributive social capital. This system is described in chapter 5. It consists of three pillars. Its first pillar is the analysis of interactions on social networking platforms and other online data sources with the help of supervised learning (chapter 6). In chapter 7, the second pillar, the use of a market system to build and maintain CSC scores, is described. As third pillar, endorsements and certifications permit recognition of CSC that people demonstrated to have in areas that are not fully covered by the other two pillars. This mechanism is discussed in chapter 8. The combination of all three pillars, the system architecture, and various design choices are discussed in chapter 9.

5 CONCEPT OF A CONTRIBUTIVE SOCIAL CAPITAL SYSTEM

5.1 GOAL OF THE SYSTEM

As discussed in section 1.1, there are several issues in online communication that range from trolling by teenagers in discussion boards to influencing whole elections. Additionally, a survey of 242 university students – described in more detail in section 1.1 – revealed that it is difficult to assess someone's trustworthiness, to identify experts, and to advertise own knowledge on OSNEM platforms. The latter would be useful to provide help to others who have questions in one's area of expertise. A hypothesis is that better transparency regarding a person's contributive social capital can mitigate these issues. As increased CSC transparency is ideally accompanied by a better association of the online persona and the person behind them, it may also drive more civil discourse (Rainie et al., 2017). Additionally, transparency and providing people with a way to reward each other for helpful contributions can be a first step to motivate people to behave in a social and supportive way towards others. This has been demonstrated in prior work that focused on incentives to promote wanted behavior (Gneezy et al., 2011).

To achieve transparency, a CSC system needs to provide CSC scores for each participant. This score is the main goal of the system and should be regarded as a measure of the person's true contributive social capital, which includes their competence, trustworthiness, and social responsibility. It should not be a pure measure of activity, therefore, it incorporates, among others, previous interactions and the feedback of others to reflect the true value a person adds to the network.

This one CSC score should paint a universal picture of a person by accumulating different measures from individual platforms. This is an important distinction to other approaches that are often limited to the platforms for which they were developed. Examples for such limited systems are the reputation management in scientometrics based on the h-index (Zhao and Li, 2009), endorsements on Linked-in, up-vote/down-vote systems on threaded discussion platforms, and to some extent the "likes" on Facebook.

Another important distinction from these other systems is the metric CSC, which describes the value added to the network by a person, as defined in section 2.1.2. Many other systems were designed with a commercial application in mind. The Klout score (Rao et al., 2015), which accumulates a single score for a user's influence, is, e.g., a tool used by marketing managers of companies to identify influencers. The CSC score, on the other hand, is a measure that is also related to altruistic behavior and mutual support and might, therefore, lead to an increase of social behavior as it creates incentives to behave in a pro-social way (Gneezy et al., 2011).

Due to the rapid increase of interactions in online social networking platforms (Lenhart et al., 2010), (Gabbriellini, 2014), (True, 2017), people have a limited time to make an assessment about their counterparts. Having a single score as output of the system meets this requirement and allows users to quickly form an opinion. However, it is also important to prevent the halo effect from happening. The halo effect describes that people can be blinded by one characteristic, e.g. physical attractiveness, and wrongfully transfer this positive judgment to other parts of their personality (Thorndike, 1920). This risk can be mitigated by providing a well-balanced score, which does not focus on a single characteristic. Further transparency can be created by providing an additional differentiation of each person's CSC score, e.g., in the form of a split by topic (see section 9.2.8).

To summarize, the main goals of the system are:

- Create transparency in online communication by providing a single CSC score for each participant that reflects, among others, their competence, trustworthiness, and social responsibility.
- Leverage existing feedback systems from social media platforms (e.g., likes on Facebook) to assess scores without the need of active user involvement. This recognizes past activity and prevents that users have to start from zero (see chapter 6).
- Include a market-like reward system that allows people to reward helpful contributions in the network and creates transparency by building CSC scores. This also incentivizes social behavior (Gneezy et al., 2011) (see chapter 7).
- Allow users to transfer offline-world assessments into the system, e.g., to recognize university diplomas that are of relevance for CSC (see section 8).
- Empower the "little people" by providing the ability to easily assess others as well as the chance to advertise themselves in the domains of their expertise (see section 16). This is important as 80.6% of participants in our survey (see section 1.1) found it difficult to appropriately advertise their expertise on social media today.
- Contextualize the CSC scores by categories to allow users to differentiate between CSC in various domains (see section 9.2.8). This prevents the halo effect from happening and allows a variety of applications (see section 16).
- Prevent or disincentivize fraudulent behavior.
- Establish contributive social capital as a counterweight to pure expertise/knowledge/influence identification, as CSC also includes social aspects like reputation and helpfulness.
- Create the scores in a transparent way to facilitate trust in the results. This goal is sometimes orthogonal to fraud prevention, as complete transparency regarding the social network analysis aspect would enable and even invite tempering.
- Implement the system in an easy-to-use way to encourage new users to join and participate.
- Enable new business models with the help of the market-based reward system and the improved transparency.

- Lay the ground for voting systems for experienced people that take a person's social capital weight into account. Thereby people who contributed to societal discussions, helped others, and were rewarded for that, would have a higher voting power. This is not directly in line with the democratic principle and therefore needs to be additionally investigated to determine whether the advantages outweigh the risks, as discussed in section 16.6.

- Facilitate the creation of a better society by introducing social thinking into the ways of the homo economicus (people are provided an incentive to contribute and help each other).

- Enable the identification of people with high CSC in different topics to ask them questions or to invite them to round tables for (political) discussions. The people with high CSC can be regarded as experts in their field. This goal directly addresses the issue in online communication that experts are hard to identify, a claim that was supported by 59.5% of participants of our survey (section 1.1).

5.2 VISION OF A CONTRIBUTIVE SOCIAL CAPITAL SYSTEM

In order to achieve the goals presented in the previous section, we envisioned a system to recreate the CSC of a person as holistically as possible. To achieve this, several of the methods to infer the CSC of others, as discussed in section 2.3, were leveraged. The result is a combination of different tools and various data sources. The system is structured along three main pillars: leveraging data from online social networking platforms, building and maintaining CSC with social capital markets, and including real-life expertise with certifications and endorsements. This is illustrated in figure 6.

The first pillar, the analysis of OSNEM, leverages information from contributions and peer feedback that is available on a variety of online platforms. This can take place without active contributions of the user and thereby facilitate the goal of easy usability and a holistic approach.

The second pillar represents a market-based reward system with its own currency. Transactions between market participants allow to recognize and reward altruistic and pro-social behavior or pay for received information or services. Simultaneously, every transaction dynamically builds up the CSC scores based on the underlying mechanism. The direct usability of the exchanged currency that can be reused goes beyond the "likes" exchanged on some social networking platforms. This allows for more incentives of social behavior and lays the ground for new business models (see section 16).

Including certifications and endorsements as third pillar are another way to recognize offline CSC within the system. Certifications of institutions like university diplomas that are of relevance for CSC can be integrated in this way. Endorsements by others have similar motivation and overlap with transactions via the market system.

All three pillars are discussed in detail in the following chapters 6, 7, and 8.

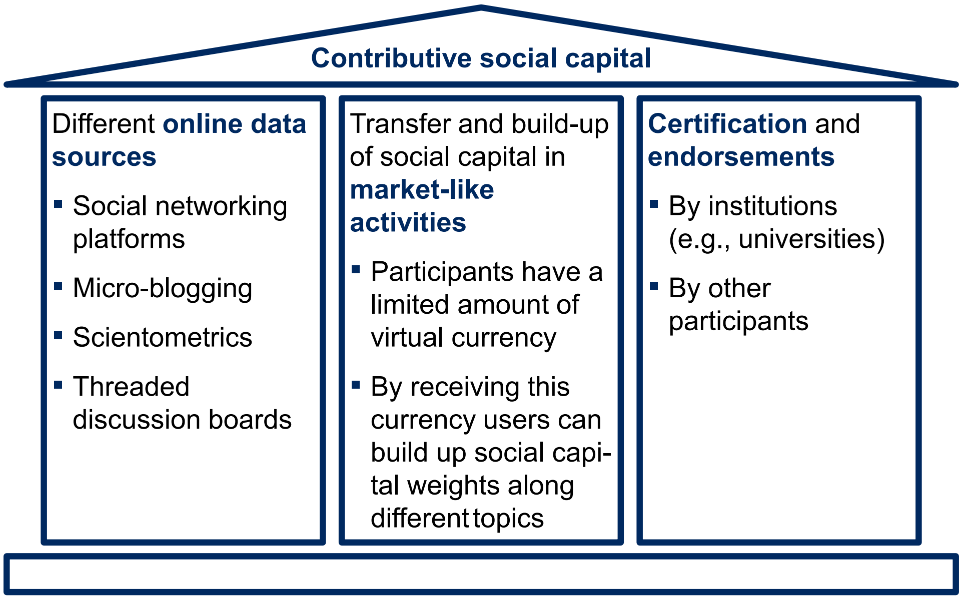

Figure 6: Vision of the CSC system along three pillars

5.3 DEFINITION OF THE CONTRIBUTIVE SOCIAL CAPITAL WEIGHT (CSCW)

The goal of the system is to measure the contributive social capital of a person. As CSC is a term that describes the whole phenomenon of contributing to the overal social capital of the network, as described in section 2.1.2, we introduce a new term as a concrete measure for each person. The **Contributive Social Capital Weight (**$CSCW_i^t$**)** of an individual i in topic t is a measure for their respective contributive social capital. This measure can be represented as a real value, e.g., on a scale from 1 to 10. In the following, this term is used analogous to CSC score or CSC value.

5.4 PRIVACY CONCERNS

As mentioned above, the system is partly build on the analysis of interaction data on OSNEM. The resulting scores condense information about the participants with the goal to make them accessible to other participants.

In the light of recent scandals about the misuse of social networking data (NYTimes, 2018), it is important to keep an eye on the privacy of the users and identify ways to prevent potential abuse of the system.

The system design (chapter 9) values individual privacy highly, e.g., by including permissions before any data is extracted. A detailed discussion about risks with regard to two current examples is offered in chapter 17.

6 SOCIAL CAPITAL EXTRACTION FROM ONLINE DATA SOURCES

6.1 MOTIVATION AND UNDERLYING IDEA

In recent years, an increasing share of social interactions took place online (Lenhart et al., 2010), (Gabbriellini, 2014), (True, 2017). There is a variety of platforms that facilitate online communication and the exchange of information. As many of the discussions and relations on these platforms are publicly available, they can be investigated for the assessment of individual CSCWs. In chapter 3, five potential data sources were introduced. All of these sources allow the extraction of user features, for instance, the number of comments posted. Some of these sources implement intrinsic feedback and voting mechanisms that visualize the support of others — which may be an additional indication of a user's social capital. The hypothesis that social capital can be analyzed on the basis of OSNEM is backed up by (Best and Krueger, 2006), who conclude from several other publications that "online social interactions meet the conditions necessary to facilitate the production of social capital".

Based on these facts, we want to investigate whether the analysis of OSNEM and other data sources can be used for the assessment of individual CSC scores.

6.2 POTENTIAL DATA SOURCES FOR THE ANALYSIS

The first step is to define relevant data sources. Contributive social capital focuses on the value a person adds to their social network in the form of knowledge or help shared via interactions with other network participants. Therefore, the most relevant data sources include interactions, rather than personal data like a person's race, sex, wealth, or whether they have a Wikipedia article. As discussed, five data sources that fulfill these requirements are social networking platforms (like Facebook), microblogging platforms (like Twitter), threaded discussion boards and Q&A portals (like Quora or Reddit), scientometrics, and direct communication (like email).

Direct communication in the form of email or Whatsapp messages is usually private. The other four sources can be, to some extent, publicly accessed and investigated for CSC assessment. A detailed discussion of all sources with respect to social capital research is provided in section 3. In general, we can note that for each user of these data sources, one can compile a list of (interaction) features. These lists are different between the sources. For Facebook they include, e.g., the number of friends, posts, and comments, as well as the "likes" received on own contributions. On Twitter some of the features are the number of followers and followees, tweets, and retweets. In threaded discussion boards one can again look, e.g., at the number of contributions and the number of up- and down-votes. Popular features in scientometrics are the number of citations and publications.

6.3 CSC ANALYSIS IN ONLINE DATA SOURCES

In order to create a CSC estimation with supervised learning, several steps are needed.

1. A data set from the source that is to be investigated is required. This data set can be created continuously by accessing the data source via an API or crawled once for an initial analysis.

2. Pre-processing of the data is necessary to remove potential double entries and create feature vectors for each user that list all relevant accessible features. It is important to include many different features to allow a CSC estimation that holistically reflects an individual's personality. Single features are sometimes used to predict individual characteristics. The number of followers on Twitter, e.g., is often used as an estimate for a person's influence (Anger and Kittl, 2011). This disregards the root cause of why the person has many followers. People regularly follow others because they find their tweets interesting or entertaining, even though they do not agree with the content or share that person's views. These metrics are also open to manipulation as followers on Twitter can be bought (Anger and Kittl, 2011). To address these shortcomings and create a CSC estimate that includes all three constituents, our analyses were tailored to the different data sources. The used features can be generally divided into four groups:

    - Personal information, like name, age, or profession.
    - Activity features that reflect how active a user is on the respective platform (e.g., number of comments, posts, etc.).
    - Feedback features that reflect the responses of other users (e.g., likes, retweets, etc.).
    - Features derived from centrality measures that denote a person's position within the network.

    The features identified on the four different online data sources are listed in detail in chapter 11.

3. To train the supervised learning algorithm, a ground truth CSC value is required. Ideally this value was compiled in a trustworthy way, e.g., via assessments from other users. Unfortunately, such an assessment was not available for some of the experiments we conducted. In these cases alternative ground truth values were investigated and used.

4. Once the algorithm is trained, it can be used to predict the CSC value of all users for whom features are available.

Based on these basic steps, we conducted several experiments that are described in chapters 11 and 12. These experiments investigate the extraction of CSCWs from different online sources. Chapter 10 summarizes previous work about the extraction of social capital and related properties from different online data sources.

In addition to these investigations, we also examined to what extent subjectivity and biases can be visualized on the same data platform. This creates transparency regarding the provenance of an individual's support (in the form of contributive social capital). These investigations are discussed in chapter 18.

6.4 contextualization by topics

The steps presented in the previous section are a way to extract a single contributive social capital weight per user based on previous interactions on contributions. In order to contextualize the CSCW by topics, an analysis of the main areas of interest, respectively the topics in which they interact, is required. Based on the hypothesis that a person's CSCW is the largest in areas in which they are most interested in, the CSCW can then be divided along those topics. Identifying areas of interest on Twitter has been successfully demonstrated with topic models like LDA (Ramage et al., 2010) and with the help of ontologies (Michelson and Macskassy, 2010).

7

BUILDING CSC SCORES WITH SOCIAL CAPITAL MARKETS

7.1 MOTIVATION AND UNDERLYING IDEA

Markets are used successfully all over the world to distribute goods and information. In capitalistic theory the market regulates the allocation of resources, as explained by Hazlitt: "If there is no profit in making an article, it is a sign that the labor and capital devoted to its production are misdirected: the value of the resources that must be used up in making the article is greater than the value of the article itself." (Hazlitt, 1979)

Using a market system for the assessment of a user's CSC is an attempt to investigate the applicability of markets as an analytical mechanism to build CSCWs. The hypothesis is that social capital markets may channel the capitalistic striving for profit towards altruistic and pro-social behavior, simply by visualizing the contributive social capital weights that reflect received payments. The possibility of receiving rewards may additionally incentivize social behavior (Gneezy et al., 2011).

The basic idea is to offer users a mechanism to provide dynamic feedback in the form of currency transactions about each other. Participants of online communities are used to providing feedback via "liking", "retweeting", or "voting" (compare chapter 3). The market mechanism goes beyond these straightforward measures, as all feedback is provided in the form of currency transactions. The currency can be reused by the recipient. This adds more value to received feedback (i.e., payments), but it also limits the amount of interactions to the available currency.

Based on the interactions, individual topic-specific CSCW scores are successively build up via market interactions and are assumed to be proportional to the user's real CSC value. This happens simultaneously thanks to a mechanism that is described in the following section.

7.2 CSC ASSESSMENT WITH THE MARKET SYSTEM

Every participant has a certain amount of virtual currency that we call social capital currency (SCC). The SCC is either distributed when registering to the system, as a monthly payment similar to basic income or coupled to other parameters. All market participants can transfer all or part of their currency to others. The reasons for such a transaction can vary. One can pay for information or services received, thank others for helpful contributions to social media platforms, or acknowledge altruistic behavior in general. However, for all payments one needs to specify a topic that relates to the part of the recipient's social capital weight that caused the transaction in the first place. The recipient's social capital weight in this topic $CSCW_{recipient}^{topic}$ is then increased in addition to them receiving the transferred currency SCC. As an example we can look at the following interaction:

Andy wants to thank Betty for helping him with his math homework. Therefore, he selects an arbitrary – unless specified differently before the interaction – amount ΔSCC of social capital currency and transfers it to Betty. As a topic he selects *math*. As a result the amount of SCC in Betty's possession is raised by ΔSCC. Also, her contributive social capital weight $CSCW_{Betty}^{Math}$, the measure for her CSC in the topic math, is increased. Betty has two advantages. She can use the additional currency to pay other participants and she can use her increased weight to advertise her math skills.

The system can be summarized with the following equation. A is the person who sends ΔSCC to person B in topic i. After the transaction A's SCC and CSCW are:

$$SCC'_A = SCC_A - \Delta SCC \tag{20}$$

$$CSCW_A^{i\prime} = CSCW_A^i \tag{21}$$

For the recipient B SCC and CSCW change:

$$SCC'_B = SCC_B + \Delta SCC \tag{22}$$

$$CSCW_B^{i\prime} = CSCW_B^i + \alpha \cdot CSCW_A^i \cdot \Delta SCC + \beta \cdot \Delta SCC + \gamma \cdot CSCW_A^i + \delta \tag{23}$$

The change for SCC is self-explanatory. For B's CSCW we have five contributing terms:

- $CSCW_B^i$ is B's CSCW before the transaction.
- $\alpha \cdot CSCW_A^i \cdot \Delta SCC$ is an increase that is influenced by a factor α, the amount of transferred capital ΔSCC and A's $CSCW_A^i$. The weight was included to achieve a PageRank like mechanism that takes the social capital weight of the senders into account, as they have already demonstrated their value to the network. The term ΔSCC weighs larger transactions (i.e. more important ones) higher than small transactions.
- $\beta \cdot \Delta SCC$ only takes the value of the transaction into account.
- $\gamma \cdot CSCW_A^i$ only takes the sender's CSCW into account.
- A constant term δ.

The CSCW of the other topics is not influenced by this transaction. In the example discussed above, β, γ, and δ were set to zero and α was set to a constant value. It is also possible to add similar terms to the calculation of SCC'_B. This was not included, as it is likely that such an addition would raise inflation in the overall system. This mechanism and the equations are discussed in more detail in section 7.3.1.

7.3 DISCUSSION OF MARKET PROPERTIES AND PARAMETERS

7.3.1 *CSCW Build-up Mechanism*

The CSCW build-up mechanism links the monetary interactions to the build-up of CSC scores as described by equations 20, 21, 22, and 23.

Equations 20 and 22 describe the monetary exchange and are similar to interactions in other markets. This is also the place where a sales or transaction tax could be introduced to fund the basic income of the system or cover other expenses. Additionally, one could implement terms that increase the received currency. This would allow an additional reward for the recipient, however, it is likely to massively increase inflation within the system.

Equation 21 reflects that the CSCW of the sender stays untouched by the transaction. This is an important concept as it makes the build-up of the recipient's CSCW "free". The average CSCW of all participants consequently increases with every transaction. We chose this straightforward mechanism in order to keep the system easy to understand and to not discourage transactions. However, if future investigations indicate the necessity, the sender's CSCW could be reduced at this point. This would prevent an inflation of the CSCWs, but it is likely to be perceived as unfair and discourage participation. Alternatively, one could introduce a normalization that keeps the average CSCW of all users constant over one topic. This would effectively – to a much smaller degree – decrease the CSCWs of everyone besides the recipient. We did not investigate this for the same reasons.

The real build-up of CSCW happens on the recipient's side, as described by equation 23. As mentioned in the previous section, the new $CSCW_B^{i'}$ is determined by five summands that describe most of the rational choices for the CSCW build-up. Let us look at different cases to understand the meaning of each term.

- $\alpha = \beta = \gamma = 0$, $\delta \neq 0$:

$$CSCW_B^{i'} = CSCW_B^i + \delta \tag{24}$$

  In this case the increase of $CSCW_B^{i'}$ is solely determined by the constant summand δ. If $\delta = 1$, every received transaction increases the recipient's CSCW by 1. The sender's weight and the transferred currency do not play any role.

- $\alpha = \beta = \delta = 0$, $\gamma \neq 0$:

$$CSCW_B^{i'} = CSCW_B^i + \gamma \cdot CSCW_A^i \tag{25}$$

  The γ-term introduces the sender's $CSCW_A^i$ in the topic i in which the transaction occurred. Thus, a transaction of a person with high CSCW increases the recipient's weight more than that of someone with low CSCW. This incorporates some ideas of the PageRank (Wu et al., 2007) and Hubs and Authority (Schütze et al., 2008) algorithms, which also determine the weight of a node within the network based on the weights of the nodes that link to it. These algorithms have been used in the past for a variety of applications that are related to social capital, e.g., for expert identification (Campbell et al., 2003), (Yang et al., 2010).

- $\alpha = \beta = \gamma = 0$, $\beta \neq 0$:

$$CSCW_B^{i'} = CSCW_B^i + \beta \cdot \Delta SCC \tag{26}$$

In this instance, the weight increase is driven by the transferred amount of currency ΔSCC. Hence, larger transactions increase the recipients weight more than small ones. The underlying reasoning is that larger transactions are signs for larger contributions by the recipient. To appropriately recognize this contribution to the network and reflect it in the weight of the contributor, the CSCW increase is weighted by the transferred currency.

- $\beta = \gamma = \delta = 0$, $\alpha \neq 0$:

$$CSCW_B^{i'} = CSCW_B^i + \alpha \cdot CSCW_A^i \cdot \Delta SCC \tag{27}$$

The α term includes elements of equation 26 and 25: The increase of $CSCW_B^{i'}$ is influenced by the sender's weight as well as the amount transferred. This combines the advantages of both terms – a PageRank like quality as well as the importance of larger transactions.

Which of the parameters α, β, γ, and δ are set to zero depends on design choices and the specific goal of the market system. For a system whose main goal it is to measure the CSC of its participants as close to reality as possible, the sender's weight and the transferred amount are certainly interesting. Therefore, equation 27 was chosen for the investigations in chapters 14 and 15. All feedback in the form of transactions is consequently weighted in two ways. First, by the weight of the sender, which assigns more influence to people with a higher weight within the system. Second, by the height of the transaction. This makes every payment act in two ways: the recipient not only receives the currency but also a proportional increase of his or her weight, which is another benefit.

The value of parameter α, and if required of the other parameters, can be defined with regard to the following two goals. On the one hand, every SCC payment should entail a visible increase of the recipient's CSCW. The reason for this is psychological, as for people to pursue goals effectively, a tracking measure that visualizes the progress is beneficial (Locke, 1996). On the other hand, the increase due to a single transaction should not be massive. This time the reason is practical, as large increases of the CSCW would quickly inflate the scale. This could be received as unfair by other participants. It is, furthermore, possible to define α in a way that it is dependent on the CSCW of the recipient: $\alpha = \alpha(CSCW_B^i)$. This allows, for instance, to create a logarithmic scale (see section 9.2.9), which was successfully used in comparable systems (Rao et al., 2015) and our market simulation (chapter 15).

7.3.2 *Currency*

"*Money* is anything that is generally accepted as payment for goods or services or in the repayment of debts." (Mishkin and Serletis, 2011) If established currencies, like the USD or EUR, are used for the social capital market system, an advantage is provided to wealthy people. As there is no reason to presume a causal relationship between wealth and high CSC, we proposed a new currency that is not linked to any real monetary systems: social capital currency. It has the properties of real currency insofar that it is – within the system – universally accepted and artificially scarce. However, there is no direct exchange to other established currencies to prevent wealthy participants from having an intrinsic advantage.

Even though SCC is not directly linked to other currency, it has an intrinsic value as currency transactions directly increase the recipient's CSCW, which can be used, e.g., as advertisement. The value of SCC can therefore be described as immaterial as opposed to the material value of EUR or USD that can be used to buy material goods. To reflect this intrinsic value, we do not use the term "play money", which is often used to describe virtual exchange tokens of no value, but social capital currency.

The use of currency that is limited instead of virtual likes that can be handed out indefinitely, has three distinct advantages. First of all, it prevents people from supporting others arbitrarily or out of reciprocity, which can be observed on other platforms, like Twitter (Weng et al., 2010). Secondly, the feeling of receiving a payment is a larger reward than a "like" and may, therefore, incentivize further social behavior (Gneezy et al., 2011). Thirdly, the currency can be used for payments, which may enable new business models and increases the realism of the CSCW due to market awareness (compare chapter 16 about use cases of the system).

7.3.3 *Distribution of Currency*

In order to allow participation within the system, a distribution algorithm to hand out currency to all users is needed.

State-controlled currency notes and coins are usually created by the department of treasury and distributed via a network of banks (U.S. Department of the Treasury, 2017). Today, even more important than notes and coins is "having credit", i.e. people and companies can pay electronically with the promise that the money will later (automatically) be transferred. As SCC is unconnected to this system another distribution method is required. There are several options to distribute SCC:

1. Distribute SCC evenly at the launch of the system,
2. hand out a defined amount of SCC whenever a new person registers,
3. SCC can be earned by providing resources to support the system,
4. exchange SCC for other currencies,
5. implement a basic-income-like distribution mechanism.

Distributing SCC at the launch of the system would encourage people to participate early on, however, users who join later would have no SCC at the start and could, therefore, not directly participate. The second option prevents this from happening but it also comes with two potential disadvantages. People could be motivated to register several accounts in order to receive the initial bonus (sybil attacks) (Douceur, 2002). More importantly, once the initial money is spent and no new SCC is earned, the user can no longer actively participate in the system. In established cryptocurrencies, the only way new currency is issued is the mining process. This encourages users to contribute computation power to the network in order to secure transactions and prevent fraud. A similar reasoning could be applied for the CSC system, as described by the third option. The disadvantage is that only people who contribute to the system – e.g., by offering computation power – would receive currency. Therefore, this distribution mechanism might be used as an addition but not as the central one. The fourth option, buying of SCC with other currencies, was already excluded in section 7.3.2 in order to not favor real-world rich individuals. The last option counters most of the disadvantages of the ones discussed so far. By continuously distributing SCC to all participants of the system, it is ensured that everyone is treated equally and has the ability to participate, which is essential to the success of the system. Potential downsides are sybil attacks and inflation due to the continuous increase of SCC via basic income. Both of which can be countered — e.g., with a thorough registration process and trade taxes. Trade taxes could reduce inflation because they reduce the available amount of currency.

Regarding the security of transactions one could leverage blockchain technology (Kosba et al., 2016) as is done for cryptocurrencies like Bitcoin or Ethereum (Tapscott and Tapscott, 2016). This aspect was investigated in chapter 19.

7.3.4 *Contextualization by Topics*

In section 5.1, we argued that a single CSC score allows to quickly assess someone's CSC and that a further division of the CSC score by topic or category may additionally provide insights regarding a person's specific areas of expertise. The latter is essential for all kinds of topic-specific interactions like providing help in specific domains.

All market transactions have to be specified by the sender regarding the recipient and the amount transferred. Additionally specifying the topic that inspired the transaction, is only a small effort and can easily be implemented in the form of a text entry or a selection from existing topics.

To efficiently map the transactions to the corresponding topics and fields of expertise, an ontology with the right degree of granularity is required. A deeper discussion of the contextualization by topic for the whole system is provided in section 9.2.8.

The theory that market systems can be used for the assessment of CSC was investigated in an experiment (see chapter 14) and a large-scale market simulation (chapter 15).

8

SOCIAL CAPITAL ESTIMATION WITH ENDORSEMENTS AND CERTIFICATIONS

8.1 MOTIVATION AND UNDERLYING IDEA

The first pillar of the contributive social capital system allows to analyze previous public interactions of participants and to infer CSCWs from them. The social capital market, the second pillar of the system, allows a dynamic and continuous CSCW assessment based on currency transactions.

In order to obtain a holistic representation of each participants' CSCW, a third mechanism may be utilized to allow the recreation of offline social capital within the system. This is especially relevant for people who have contributive social capital, i.e. are competent, trustworthy, and socially responsible, but who are not active on social media. Their participation in the social capital market will over time depict their CSCW, however, this may be accelerated by allowing a direct transfer in the form of endorsements and certifications.

According to the Oxford Dictionary for English, endorsements are "the action of endorsing someone or something" (Stevenson, 2010) and certifications are defined as "the action or process of providing someone or something with an official document testing to a status or level of achievement." (Stevenson, 2010) In the context of social media, it has been demonstrated that personal endorsements can influence the decision process of network participants (Messing and Westwood, 2014), which underlines the importance of endorsements for the CSC system.

In this thesis we mainly focus on the first two pillars of the system. Additionally, the following sections provide a brief discussion of the endorsement and certification pillar.

8.2 CERTIFICATIONS

In the context of our system, certifications describe the process of including achievements and competence assessments from everyday life into the system. There is a variety of institutions that provide certifications that are of relevance for social capital. Following the definition of CSC in section 2.1.2, we can identify different sources of certifications for the three main components of CSC:

- Factual competence can be certified by educational institutions. All diplomas, scientific degrees, and language certifications reflect and confirm that a person has dedicated a certain amount of time to acquire knowledge or a skill that increased their competence. Certifications that confirm experience in a certain field, e.g., employments references, also fall into this category.

- Trustworthiness may also be certified by institutions, e.g., in the form of clearance certificates by the police or character certificates by previous employers or charitable organizations.
- Certifications for social responsibility may include references for voluntary support of charitable or governmental institutions.

All of these certificates reflect activities that are of relevance for a person's CSC but are not necessarily reflected in the assessments of pillars 1 and 2.

8.3 ENDORSEMENTS

Endorsements work on a personal level. Instead of institutions that issue the certifications, it is private people who attest that individuals have displayed achievements of relevance for social capital. In this regard, the endorsement overlaps with the market system, where each transaction can be seen as a micro endorsement of the recipient by the buyer. Endorsements by experts are in general regarded more trustworthy than endorsements by others (Biswas et al., 2006). This can be taken into account by considering the social capital of the person issuing the endorsement. This is similar to the way the CSCW of the sender is taken into account when calculating the CSCW increase of the recipient in the market system (compare equation 23).

8.4 DETERMINING THE AMOUNT OF THE CSCW INCREASE

In the previous sections we argued that using endorsements and certifications can be a suitable way to reflect existing contributive social capital within the system. The increase of the CSCW as a consequence of the certifications and endorsements needs to be defined in a universal way. One could define a constant CSCW increase for each certificate or endorsement, but ideally the CSCW should increase depending on two main factors:

- The amount of time required to obtain the skill or knowledge that is being certified
- and the CSCW of the institution providing the certificate or the person issuing the endorsement.

Similar to equation 23, which describes the CSCW increase due to market interactions, we can define an equation that describes the new $CSCW_B^{i'}$ of person B in topic i after an endorsement or certification from person or institution A:

$$CSCW_B^{i'} = CSCW_B^i + \epsilon \cdot CSCW_A^i \cdot \Delta T + \zeta \cdot \Delta T + \eta \cdot CSCW_A^i + \theta \tag{28}$$

In this case $CSCW_A^i$ is the weight of the person or institution certifying person B's CSCW. ΔT is a measure for the time required to achieve the skill or knowledge that is being certified/endorsed. For a bachelor's degree that requires on average to work for eight hours a day on 200 days for each of three years, ΔT would be in the area of 4,800 hours. Online certificates or other education and experience can be credited in a similar way. The duration ΔT of different certificates and endorsements need to be defined before enrollment of the system, in order to ensure constant ratings. A person should, e.g., not receive a larger/smaller CSCW increase if he or she took longer/less than 4,800 hours to obtain the bachelor's degree. What is important is the average amount of time required for the achievement.

The different summands of equation 28 describe different ways to define $CSCW_B^{i'}$:

- $\epsilon = \zeta = \eta = 0, \theta \neq 0$:

  $$CSCW_B^{i'} = CSCW_B^i + \theta \tag{29}$$

  In this case, every endorsement or certification would increase $CSCW_B^{i'}$ by a (constant) term θ. This negates the fact that endorsements and certifications may weigh differently depending on the person or institution issuing the endorsement or certification. It also does not distinguish between different certificates (a bachelor's degree would have the same impact as much as a brief online course).

- $\epsilon = \zeta = \theta = 0, \eta \neq 0$:

  $$CSCW_B^{i'} = CSCW_B^i + \eta \cdot CSCW_A^i \tag{30}$$

  In this instance, $CSCW_B^i$ is increased by a term that depends on the weight of the issuer. This is an improvement over the previous equation, however there is still no distinction between different types of certificates. A bachelor's and a master's degree would, therefore, result in the same CSCW increase, as long as the same university issued them.

- $\epsilon = \eta = \theta = 0, \zeta \neq 0$:

  $$CSCW_B^{i'} = CSCW_B^i + \zeta \cdot \Delta T \tag{31}$$

  Thanks to the term ΔT, the time required to achieve the endorsed skill is taken into account. But an endorsement from an average person would still be as valuable as that of a Nobel laureate. ζ maps the time required to obtain the skill to the corresponding social capital.

- $\eta = \zeta = \theta = 0, \epsilon \neq 0$:

  $$CSCW_B^{i'} = CSCW_B^i + \epsilon \cdot CSCW_A^i \cdot \Delta T \tag{32}$$

  This equation incorporates both, the CSCW of the issuer and the time required to achieve the skill.

The latter case, $\eta = \zeta = \theta = 0$ and $\epsilon \neq 0$, includes both time ΔT and $CSCW_A^i$ and is, therefore, our suggestion for an endorsement and certification mechanism. Alternatively, other combinations, e.g., $\epsilon = \theta = 0$ and $\zeta \neq 0$ and $\eta \neq 0$ may also satisfy the requirements. The ideal setup may be determined in field studies, especially with regard to the values of the parameters.

The parameters ϵ, ζ, η, and θ allow to fine-tune the system. This is important as endorsements and certifications should not disbalance the other assessments.

8.5 POTENTIAL CHALLENGES

The system should reflect the true CSC of a person. As discussed previously, the CSC depends on different aspects, including the competence of a person and their willingness to share and help others. By introducing endorsements and certifications in a manner described above, it is certainly possible to account for some of the aspects of a person's CSC and visualize it within the system. However, it is possible that a double counting takes places. Double counting can be best explained at the hands of an example. The three people A, B, and C are all registered to the CSC system. A and B both hold bachelor's degrees in computer science and are equally educated. B and C are both active members on an online social networking platform on which users help each other with questions about computer science. As a result B and C receive regular payments in this field and thereby increase their weights. Under the assumption that B and C display the same kind of engagement and provide similar help, their weights are equally high. After certifying their bachelor's degrees, A's and B's weights are increased. There are two consequences that might lead to misunderstandings:

- B's weight is now higher than that of person C, even though their support on the one platform they both engage on is the same.
- In this example, person A did not help anyone with any computer science related problems. Nevertheless, the weight is increased because of A's competence demonstrated by the degree.

Both points might demotivate person B. However, the increase in weights accounts for real competence that might prove useful in the future.

There is another potential problem that might occur if person B had certified the degree directly when registering to the system. It is possible that people would then transfer less SCC to B than to C because they reflect the same level of knowledge and B is – due to the certification – already ranked higher. This would be unwanted as the timing of a certification should not matter. If this bias truly takes place needs to be investigated in future work and appropriate counteraction taken, if necessary.

A third issue that may arise when allowing all participants to endorse each other's CSC, may be that of fraudulent behavior. In order to provoke a CSCW increase in the market system, one needs to transfer currency, which is limited. The amount of certifications is theoretically not limited, which may incite people to endorse others beyond their real CSC score. This can be prevented or at least limited by clear endorsement procedures and transparency. Clear procedures that, e.g., require a detailed description of why the person should receive a CSCW increase, expand the time required to complete an endorsement and thereby reduce the likelihood of people misusing this method. Transparency regarding who endorsed whom creates a feeling of accountability, which has been demonstrated in the context of organizational records management to prevent fraud and corruption (Palmer, 2000).

In all cases, a detailed process for all endorsements and certifications is required. While we listed most of the requirements in this chapter, additional investigations, e.g., regarding the parameters listed in section 8.4, are required before this process can be tackled.

9 SYSTEM ARCHITECTURE

The last four chapters presented a vision of a contributive social capital system that may help network participants to assess each other's CSCWs and encourage social behavior.

In this chapter, we have a look at how the three pillars can be combined and investigate how specific design choices could look like. The different options are discussed and a suggestion is provided. Even though we ran several experiments with a couple of hundred participants, we could not investigate the behavior of the complete system but had to focus on the parts of the system that are most interesting from a computer science perspective. Thus, not all of the design choices were verified in an experimental setting. This chapter closes with a look at potential challenges and shortcomings before parts III and IV present comprehensive experiments to investigate pillars 1 and 2 of the CSC system. In the discussion of these experiments many of the design choices are revisited.

9.1 COMBINING THE THREE PILLARS

The goal of the system is a holistic CSCW assessment for each participant that is build on input from the three different pillars. To understand how these three inputs can be combined, we briefly review how they are used to create contributive social capital weights for each user.

The analysis of OSNEM and other online data sources, as discussed in chapter 6, reviews the entirety of a user's online presence — at least to the extent for which permissions were granted. The interaction features from different social media platforms and other sources are collected and a $CSCW_{\mathtt{pillar1}}$ is calculated with the help of supervised learning.

The market system (chapter 7) allows for a dynamic assessment of CSCW scores and requires active input by the users. By specifying transactions that can be used as payments for information or services, or simply to appreciate helpful or social behavior, individual $CSCW_{\mathtt{pillar2}}$ scores are formed.

Endorsements and certifications, as described in chapter 8, are a way to acknowledge CSC that an individual has demonstrated to have in the offline world within the system. Each certificate that is registered in the system and each endorsement increase a user's $CSCW_{\mathtt{pillar3}}$.

The straightforward merging point of the three pillars are the resulting weights, $CSCW_{\mathtt{pillar1}}$, $CSCW_{\mathtt{pillar2}}$, and $CSCW_{\mathtt{pillar3}}$.

Even though it may be interesting to make each weight individually accessible to the users, only a combination of the three weights can achieve the goal of providing a single metric that is quickly comprehensible. The combined $CSCW_A$ of person A can consequently be defined as:

$$CSCW_A = CSCW_{A,\mathtt{pillar1}} + CSCW_{A,\mathtt{pillar2}} + CSCW_{A,\mathtt{pillar3}} \tag{33}$$

There is no need for individual weights when adding the different pillars, as the CSCWs can be scaled individually when they are calculated. For $CSCW_{A,pillar2}$ and $CSCW_{A,pillar3}$ this can happen in equations 27 and 32 via the parameters α and ϵ. $CSCW_{A,pillar1}$ does not have an individual parameter for its adjustment, however, its scale can be changed by adjusting the ground truth values when training the supervised machine learning algorithms. The scaling of the resulting score, is an important design choice, which is discussed in section 9.2.9.

9.2 DESIGN CHOICES

This section discusses specific design choices and different options regarding the implementation of the system. Beginning with the registration of users, we discuss the question of whether the system should be set up in a central or distributed way, and how the three pillars can be implemented individually. This chapter is rounded off with discussions regarding the division by categories, scaling options, and potential shortcomings.

9.2.1 *Registration to the System*

Each participant needs to have a unique user profile to

- grant the required permissions for the analysis of their different OSNEM profiles,
- receive SCC in the form of basic income,
- allow others to send them currency,
- upload certificates from institutions,
- provide endorsements about other participants,
- receive endorsements from others.

It is important that each participant has only one unique profile so he or she can be identified by others and not wrongfully receive multiple basic incomes. Douceur concluded that a central entity is required to assign unique IDs to prevent such sybil attacks from happening (Douceur, 2002). Dinger and Hartenstein demonstrated that this might also be achieved in distributed networks with an identity registration procedure called "self-registration" (Dinger and Hartenstein, 2006).

9.2.2 *Central vs. Distributed Setup*

An important design question is whether the system should be set up in a central or a distributed way. Both options have different advantages. The advantage of a distributed system where everything is distributed and all participants have similar rights, is that there is no central institution that may be corrupted. This ensures that there is no misuse of the system from the side of a central institution that may sell user data or push their own agenda. Consequently, the user's trust towards the system is high, which is essential for a CSC assessment system to work.

On the other hand, there is a variety of system features that are difficult to realize in a distributed environment. While sybil attacks in the registration process may also be countered in a distributed setup and blockchain technology provides a means for safe transactions (Nakamoto, 2008), there is a need for a central platform due to other reasons. The OSNEM analysis algorithm needs to be maintained and not made publicly available in its entirety to prevent tempering of the scores. The certifications and endorsements need to be entered and reviewed. Finally, the market system and basic income need to be administered and adjusted in the case of new issues – e.g., introduce taxes when the perceived inflation becomes too large. All of this can also be achieved in a decentralized way but it is much easier to implement with the help of a central institution.

While we like the idea of a distributed system and still discuss options of how to realize some design choices in a distributed environment, the above arguments point to a centralized solution.

The main issue with a central institution managing the system is potential mistrust towards their intentions from the side of the participants. An example for misuse is the harvesting of Facebook user data by Cambridge Analytica in 2016 (NY-Times, 2018). This may be countered by selecting a trustworthy institution for the administration. Private scientific institutions, like the Nobel Committee, may be a suitable choice, as long as they are independent. Governmental institutions that are managed by one country or ideally by a coalition of several countries, like the United Nation, may be another appropriate and trustworthy choice. In both cases open structures and communication are required to foster the development of trust.

9.2.3 *Online Platform*

A central web portal is required as interface for participants to use the functions of the system of all three pillars.

For the first pillar, participants need to grant permission to access the APIs of the social networking platforms and other sources they want analyzed for their CSC assessment.

The same web interface can be used to receive SCC, make transactions, and check the current SCC balance, all of which is important for the second pillar.

For the third pillar, the functionality to upload certifications and provide endorsements is needed.

Finally, the platform allows to review the own CSCW status and to make it accessible to others, in order to promote own expertise. This can also be achieved by embedding the system into existing social networking, microblogging, or Q&A platforms, as described in section 16.1. Participants can use these platforms to pose questions to people with a high CSCW in the topic of the question. A useful answer can be acknowledged via a SCC transaction in the market, which builds the CSCW of the person who provided the answer.

9.2.4 *CSC Extraction from Online Data Sources*

There are three more design choices that have to be discussed regarding the first pillar.

The first is the frequency of the analysis. There is a variety of options that range from a one-time assessment to a continuous analysis. The single assessment, either when the system is rolled out or at the point in time someone registers, has the distinct advantage that it is inflexible and soon outdated. The second and third pillar allow for a further development of CSCW scores but they require active participation on the side of others. A continuous and ongoing analysis of social networking platforms is computation-time consuming but allows for an up-to date assessment of each participant's CSCW. It has been demonstrated by Rao et al. that such an analysis is feasible even on a large scale (Rao et al., 2015). A potential downside of a continuously updating CSC score is that it may decrease relative to other users after periods of inactivity. This needs to be communicated in a transparent way to prevent confusing the users.

The second design choice regards the usability of the CSCWs in the context of OSNEM: whether or not the resulting CSCW should be embedded and displayed on social networking platforms. The weight could, e.g., be displayed next to the user name, or when hovering over it with the cursor. The advantage is apparent — it would allow participants to asses their counterpart without the detour of a different website. This would greatly improve the range of the CSC system. The choice to do this has to be made by the providers of the OSN or other social media portals.

The third design choice is if and how the input of different social networking platforms should be combined. Not combining the scores would allow a platform-specific CSCW. The downside is that people with high CSCW from other platforms may not be found as reliable as they should be, which would be a loss. Therefore, and to allow an assessment that is as holistic as possible, we suggest the calculation of a single score based on features collected from all platforms.

Depending on the nature of the platform, the extracted features can vary. A scientometrics data set, e.g., does not include follower relationships, but co-authorships or citations. A detailed discussion of the available features that we used for the analyses is provided in chapter 12. Once a common data set is created, it can be used for the training of the supervised learning algorithm. The score created in this way is holistic, can directly be assessed by the users, the computational time is not artificially inflated, and CSC that was demonstrated on one platform can be carried over to others.

An alternative is to compute an individual score for each platform and then combine the respective scores to create a universal CSCW assessment. This approach allows to assign a specific weight to the various data sources, or to combine them in a hierarchical way, which was done by Rao et al. to compute a single Klout score for each participant (Rao et al., 2015). The resulting score is also a holistic assessment of the person. We could not investigate a combination of the different CSC assessments in our experiments in chapter 11, because there was no overlap of users in the different data sets.

9.2.5 *CSC Assessment with Social Capital Markets*

In the preceding discussions in this chapter and in chapter 7, the most important design choices regarding the market system were already covered. People receive SCC on a centralized platform, can review their balance and send SCC to other participants in order to pay for information or services, or to acknowledge helpful and pro-social contributions. An important aspect is the security of the system, which may be warranted by leveraging blockchain technology. Our investigations regarding an implementation on the blockchain are summarized in chapter 19.

Two other aspects to keep in mind are taxes and fraud prevention.

In the market experiments (chapter 14 and 15) we experienced an increase of the average transaction price over time, which can be understood as inflation. The reason for this increase is likely the nature of the implemented basic income. After every time period new SCC is created and distributed to all participants. This increases the overall available currency and the participants react by spending more for the same goods. One lever to counter these effects are taxes. Taxes could come in several flavors, e.g., trade (sales) taxes, income taxes, or capital taxes. All of these charges reduce the overall SCC available to participants and thereby the risk of inflation. The collected SCC could also be used to pay people for supporting the CSC system in a fashion similar to the mining of cryptocurrency.

The increase of the recipient's CSCW with every transaction paves the way for a new type of fraud. By transferring a defined amount of SCC back and forth between two users, they could push each other's CSCWs. We call this unwanted mechanism *pumping* (of CSCW). We did not experience this behavior during the market study. If required, modern fraud prevention techniques, like the ones used to detect credit card fraud from transaction and network features (Van Vlasselaer et al., 2015), can be leveraged to identify and eliminate this behavior in the CSC market.

9.2.6 *CSC Assessment with Endorsements and Certifications*

Certifications, e.g., university degrees, can be uploaded to the central platform and then assessed regarding the type of degree, the issuing university, and other criteria, as described in chapter 8. This assessment can be conducted or supported by automated systems that are similar to those used for recruitment in companies and have demonstrated to perform consistently compared to human recruiters (Faliagka et al., 2012).

To prevent fraud and foster transparency, the uploaded certificates and endorsements can be made public on this platform. This allows others to validate the findings and report suspicious behavior.

9.2.7 *Policing and Sanctioning*

The possibilities of fraud and misuse within the system range from pumping to fraudulent endorsements. It is also possible that people try to buy SCC with EUR or USD in order to increase their CSCWs and use them for advertisements. Fraud detection algorithms can be used to automatically detect and flag such fraudulent behavior

(Van Vlasselaer et al., 2015), as discussed in the previous sections. This leads to the question of how to proceed, if such cases are detected.

Two main goals of the system are to improve online communication and to encourage social behavior. These goals are altruistic and create value for every participant. Being sanctioned by being prohibited from using the system for a while, or even permanently may, therefore, be a suitable deterrent.

Online discussion portals, like Slashdot or Reddit, often employ volunteers as moderators that review the interactions and sanction participants if necessary. A similar approach may also limit fraud in the CSC system.

9.2.8 *Contextualization of CSCW by Categories*

The contextualization by categories is an essential part to unfold additional options for the usability of the CSC scores.

Without a split along different categories, the CSC value is still a measure for the overall contributive social capital of a person and can be used to identify important members of the community. Having the score available along different categories, allows people to identify the individuals with the highest CSCW in different topics of interest, which is especially helpful in Q&A portals or for related interactions. There are two main options regarding the type of categories:

- Contextualize CSC by topics.
- Contextualize CSC by topics and additional, orthogonal characteristics.

Before these two options are discussed, we briefly review the concept of ontologies that may be leveraged for the contextualization.

9.2.8.1 *Ontology Engineering*

Following Studer et al.: "an ontology is a formal, explicit specification of a shared conceptualization." (Studer et al., 1998)

In this context, a *conceptualization* can be described as tuple (D,R), where D represents all elements in the system (e.g., all topics relevant for contributive social capital), and R all relations between the entities in D (e.g., "is a subtopic of", "is the same as"). With the help of (D,R) and suitable axioms, the ontology can be defined in a formal machine readable way. (Guarino et al., 2009)

Once an ontology is formalized, it can be used to "share common understanding of the structure of information among people or software agents." (Noy and McGuinness, 2001) This underlines that ontologies are usually not the goal by themselves but simply tools with various use-cases (e.g., the aggregation of knowledge).

According to Iqbal et al. "an ontology engineering methodology caters the methodological aspect of ontology development. It gives a set of guidelines and activities to develop ontologies." (Iqbal et al., 2013)

For the CSC system we are interested in ways to develop an ontology to create a common understanding of the topics and their relations with each other. This can be achieved with methodologies, e.g., the NeOn Toolkit[1]. The RDF (Resource De-

1 http://neon-toolkit.org/wiki/Main_Page.html (retrieved: 2018-04-18)

scription Framework), which was introduced by the W3C[2] is also helpful to describe relations in a machine readable way.

A comprehensive review of methodologies and tools for building ontologies is offered by (Corcho et al., 2003) and (Gómez-Pérez and Benjamins, 1999).

9.2.8.2 *Contextualization by Topics*

The contextualization by topics is similar to the structuring of many online communities that provide thematic categories. In the context of CSCW, it allows users to identify experts with high trustworthiness and social responsibility in topics that are of interest to them.

The direct approach to achieve a thematic split, is to allow each user to individually tag contributions with topics. This has the advantage that it requires little effort from the system provider and allows people to use the topics they want. The problem with this approach, however, is that it can create ambiguities. A person may, e.g., receive SCC in the topics "math" and "mathematics", which should build up the same CSCW instead of two different weights. To counter this, a topical structure, e.g., with the help of an ontology is required.

There are several ways to provide such a structure.

The easiest way is to have a selection of topics that are all on the same level and that are collectively exhaustive and mutually exclusive. That means that all possible topics are covered by the selected "buckets" and that there is no ambiguity when users assign a contribution to one of the them. The questions to be answered for this split are, what the most relevant topics are and how many of them should be selected. These investigations may be guided by the most popular topics on Twitter that were investigated by Petrović et al. (Petrović et al., 2010) The advantage of this approach is that it is easy to understand and use. However, some nuances may be lost because the topics are on an overarching level. Some ambiguity may also still exist, as not all contributions can be uniquely assigned to only one topic. The act of teaching someone how to solve differential equations may, e.g., be attributed to either "teaching" or "mathematics".

Dividing the topics in an hierarchical way accounts for subcategories. This can be achieved by organizing the topics in a graph, similar to the hierarchical structuring of communities in social media described by Papadopoulos et al. (Papadopoulos et al., 2012) In this way, topics that are related, like "physics" and "chemistry" can both be attributed to the category "science".

Even more flexibility is provided by organizing the topics in a graph. In this way, all related topics can be connected and topics can be assigned to multiple categories. An example for such a graph structure that is organized as an ontology (see section 9.2.8.1) is DBpedia, which is extracted from Wikipedia (Mendes et al., 2012).

A fourth option is to use the topic distribution θ_d of every contribution to assign a CSC weight to all corresponding topics in a "fuzzy" way (see also topic modeling, section 4.4).

For our experiments we chose the first option, as no great complexity was required for the relatively small data sets. This may change when creating a larger CSC system.

2 https://www.w3.org/TR/WD-rdf-syntax-971002/ (retrieved: 2018-04-18)

If one of the other three options are realized at a later point in time, two properties should be included. First of all, the participants should be able to collectively adjust the topics in order to account for new subjects, e.g., with the help of the initially stated tagging of topics. Secondly, the CSC someone has in one topic should propagate along the branches of a tree and the edges of a graph. I.e., an increase of the CSCW in, e.g., the topic "physics" should consequently increase the person's weight in the topic "science".

9.2.8.3 *Contextualization by Topics and Characteristics*

Theoretically there is no need to stop at the division along thematic categories. Users could be interested in the subcomponents of CSC or other related properties, like influence or reputation. However, opening the system to this possibility brings two distinct disadvantages:

- The readability is decreased, when additional categories – maybe even per topic – are introduced.
- As the system should be open to change through the community it could lead to characteristics as categories that not all participants agree with. These characteristics could be religious or political preferences, skin color, attractiveness, or sexual orientation.

To avoid misunderstandings this thesis is limited to the research of a division by topics. However, all investigations can be extended to other categories as well.

9.2.9 *Scaling and Visualization of CSCW*

Another important design question is the representation of the CSCW within the system. Many different scales and rating schemes are used for all kinds of different applications. Letter-based ratings are used to assess the creditworthiness of countries and companies from AAA to D (Everling, 2013). Easy numerical scores are used to assess academic performances in schools and universities. Scientific applications also often leverage numerical scales, usually with finer steps. To depict large ranges in a straightforward way, a logarithmic representation may be selected. This is the case for the decibel and Richter scale.

For the representation of CSCW in the system, the scale needs to be able to reflect fine nuances that are the result of small SCC transactions. This cannot be achieved with a letter-based or an integer scale but with linear or logarithmic scales with real numbers. In order to decide which representation is the best for the success of the system, psychological considerations are helpful.

Any representation of CSCWs allow people to compare their own performance with others. Comparisons with people who are better off, e.g., regarding their performance have been extensively studied in psychological research (Garcia et al., 2013). Such upwards comparisons, which are bound to happen when visualizing CSC scores in an easy-to-grasp manner, have been associated with negative effects and can be "ego-threatening" (Neighbors and Knee, 2003). Other studies have demonstrated that, in some contexts people prefer to interact with others who are in a similar position or only slightly better off (Molleman et al., 1986).

Based on these investigations, a scale may be preferable that does not exaggerate the differences between the CSCWs of people but still honestly depicts different weights. For a CSC system that is used over the course of several years, with many participants and transactions, this is only fulfilled with a capped scale. A linear scale may result in differences in the size of several magnitudes. A logarithmic scale that ranges, e.g., from 1 to 10 may, therefore, be ideal. In short term experiments, where the CSCW of individuals is not expected to increase rapidly, a linear scale may also be sufficient. The lower cap at 1 instead of 0 allows to calculate the influence of weights in a direct way (compare equation 27) and prevents negative connotations that may arise if a new participant is associated with a 0 in the system.

9.3 POTENTIAL CHALLENGES AND SHORTCOMINGS

This section describes potential risks and shortcomings regarding the design of the system. Shortcomings of the experiments are discussed in the respective chapters. In chapter 21, potentials and limitations are discussed for the whole system.

One of the major risks of a CSC system are **privacy concerns**. To successfully infer CSCWs from interactions in OSNEM, vast transparency regarding these interactions is required. Additionally, the uploading of certifications and the need to register with an ID to prevent sybil attacks, requires further transparency from the side of the user. Even though the goal of the system is to use this data for benevolent purposes, the need for transparency may be regarded as major downside of the system. As discussed above, this may be countered by a trustworthy committee in charge of the administration. There is also previous work about preserving privacy in online communities and clouds, especially in the context of patient's health data, that may be useful when rolling out the system (Abbas and Khan, 2014). Additionally, every participant should be able to decide individually which permissions he or she wants to grant. The CSCW is then only calculated based on the resources made available by the users. The visibility of the score can also be restricted, similar to the privacy settings known from Facebook. This prevents unsolicited inquiries to people with high CSCW.

A related issue is the possibility of **misuse of the system by central organizations** similar to what we have seen on Facebook during the 2016 presidential election (NYTimes, 2018). The Chinese social credit system (see chapter 17) is an example for a whole system that aims to assesses the overall "social credit" of citizens and has been criticized for the punishments of people who do not behave in the way deemed appropriate by the government. Both issues can be prevented by a distributed setup or a well-trusted central international organization, as discussed in section 9.2.2 and chapter 17.

The system was designed under the assumption that people are generally benevolent and do not try to cheat each other. When rolling out the system on a larger scale, we need to account for fraudulent behavior. **Pumping CSCW** and **false endorsements** are intrinsically possible within the current design of the market and endorsement system. Both require the use of protective algorithms or moderators and appropriate sanctioning, as described above.

Another form of fraud are **sybil attacks**, as described in section 9.2.1. They can likely be prevented by establishing a thorough registration process. Alternatively there are ways to detect multiple accounts by one user with the help of Graph-based Sybil Detection (GSD) algorithms that were investigated in the context of social networks and P2P networks (Boshmaf et al., 2013), (Danezis and Mittal, 2009).

These protective algorithms, as well as the analysis of OSNEM activity require a large amount of **computational power**. The calculation of the Klout score, an assessment of a person's influence in the form of one single score from OSNEM activity, requires to read 18.5 TB of new data each day (Rao et al., 2015).

Part III

INVESTIGATIONS OF CSC EXTRACTION FROM ONLINE DATA SOURCES

The following chapters describe the experiments carried out to investigate the first pillar of the CSC system: the extraction of contributive social capital from online data sources. They are designed to cover all relevant aspects of CSC assessment. In chapter 10, previous work about the analysis of properties related to CSC is reviewed. The investigated data sources are social networking platforms, microblogging, direct communication, scientometics, and threaded discussion boards. Chapter 11 describes an experiment we conducted with users who provided ground truth assessments and participated on a dedicated social networking platform. Chapter 12 expands these investigations with analyses of the real-life platforms Facebook, Twitter, and Quora and a scientometrics data set.

10

PREVIOUS WORK ON THE EXTRACTION FROM ONLINE DATA SOURCES

Five popular data sources for CSC extraction were introduced in section 3. In this chapter, previous work about extraction mechanisms of characteristics related to contributive social capital are discussed and set in relation to CSC. A brief summary is provided at the end of each section.

10.1 MICROBLOGGING PLATFORMS

Twitter has been studied in a large number of research papers. However, there is no consensus on what exactly social capital is on Twitter and how it is measured. The main research focus often lies in social influence, especially identifying influential users. This is a question relevant not only for scientific research but also for the marketing departments of companies. Handing out free samples to the most important influencers can lead to successful advertising because these influencers often share their positive experiences and can spread information effectively throughout the network (Rao et al., 2015). There are several commercial services that analyze influence on Twitter. A prominent example is the Klout index, which is discussed in section 10.2. The nature of discussions on Twitter is content-oriented (Robles, 2011). Therefore, it is reasonable to assume that influential people on Twitter also have a higher than average CSC. With the help of the intrinsic metrics provided by Twitter and described in section 3.1, one can easily define more detailed performance indicators.

10.1.1 *Measuring Influence with Performance Ratios*

In their 2011 paper Anger et al. (Anger and Kittl, 2011) described and analyzed several performance ratios that are relevant for measuring someone's influence. The goal of each of the ratios is slightly different and described in the following. The same abbreviations as in (Anger and Kittl, 2011) are used.

- The *Follower/Following Ratio* (r_f) compares the number of users who have subscribed to the updates of user A with the number of users that user A is following (Anger and Kittl, 2011). The larger this number is, the more people are following the user without a reciprocal relationship.
- To detect how many of user A's tweets imply a reaction from the audience, the *Retweet and Mention Ratio* (r_{RT}) can be determined. This is the number of tweets that are amplified or lead to a communicative action between user A and another user, divided by the total number of tweets of user A (Anger and Kittl, 2011). This is a direct measure of influence of one's tweets. The more reactions a post provokes, the more influence it has had on other network participants.
- To assess how many different individual users interact with user A, the *Interactor Ratio* (r_i) can be calculated. This is the number of individual users who

retweet content or mention user A, divided by the total number of followers of user A. (Anger and Kittl, 2011)

- The *social networking potential* (SNP) was introduced to take the content focus of the Retweet and Mention Ratio as well as the relational focus of the Interactor Ratio into account. The arithmetic mean of r_{RT} and r_i forms this new metric. The SNP thereby combines many important features: number of followers, individual interactors, retweets, mentions, and tweets. (Anger and Kittl, 2011)

Anger et al. analyzed these ratios on a data set of Austrian Twitter users from 2011. Their findings, with regard to the four different ratios, are summarized in the following, and the advantages and shortcomings of each method are discussed.

- The number of followers is one of the most often used metrics for influence on Twitter. By itself this number is prone to error: many people follow others in the pure hopes of being followed back — which is often successful (Anger and Kittl, 2011). Furthermore, there are several services that sell followers (Anger and Kittl, 2011). In addition to that, the pure number of followers does not tell anything about their own activity level or social capital. All these issues can skew the effectiveness of the r_f. Anger et al. created a list of the top ten Austrian Twitter users, ranked by number of Austrian followers. The number one ranked person has a r_f of 130.1, meaning that the number of followers is about 130 times larger than the number of people he is following. Most of the other users in the top ten list, however, have a single-digit score. One user in the top ten list has a ratio of only 1.0 — he is following as many users as are following him. Considering that he is still among the user's with the most followers illustrates the shortcomings of the Follower/Following Ratio.

- The Retweet and Mention Ratio is a measure of how involved the influencers are with other users. The top three Austrian Twitter users identified with this measure are different to the top three found with the r_f metric. This illustrates the different nature of both ratios. (Anger and Kittl, 2011)

- Ranking the Austrian Twitter users according to the Interactor Ratio results in another top three. The differences are not surprising, because this ratio focuses more on conversation-oriented interactions, as it is increased by every mention or retweet a person receives. (Anger and Kittl, 2011)

- The SNP is of interest to us because it takes content (r_{RT}) as well as personal relations (r_i) into account and both are important for social capital, which is influenced by information access as well as relations. The resulting top three are all known from the top three lists of the previous ratios, which is due to the fact that the equation is the arithmetic mean of r_i and r_{RT}. (Anger and Kittl, 2011)

Especially the content-focused r_{RT}, as well as the SNP and to a smaller extent the r_i may be of interest for social capital research. People score high on these ratios when their post quality is high (r_{RT}, SNP), which reflects that they contribute to the network, or that they have sound relations with others (r_I), which is also an indicator for value within the network. High values on these ratios may consequently be indicators that the person has a high contributive social capital. Nevertheless, there are also examples of people who would score high on these ratios but who add little value to the social capital of their respective networks. A CSC prediction consequently needs to regard more features to achieve a reasonable and holistic assessment, as discussed in section 6.3.

10.1.2 *PageRank Algorithm for Social Network Analysis*

Other work employs and extends Google's famous PageRank algorithm (see section 4.1.2.5).

A similar notion can be applied to users of social networks instead of web pages. In this case the nodes are users instead of websites and the edges are votes exchanged between users, or similar metrics, instead of hyperlinks. Instead of the importance of a website, one can then analyze the contributive social capital of a user within the network. An early attempt to use a PageRank-like algorithm to research Twitter was the introduction of TunkRank by Daniel Tunke[1]. It is defined with a recursive equation that calculates the influence of a person X based on the influence of the followers:

$$\mathrm{Influence}(X) = \sum_{Y \in \mathrm{Followers}(X)} \frac{1 + p \cdot \mathrm{Influence}(Y)}{||\mathrm{Followers}(Y)||} \qquad (34)$$

Influence(X) is the estimated number of people who will read a tweet by person X. The probability that person Y retweets person X is assumed to be the constant factor p and the probability that Y reads a tweet by X is assumed to be $\frac{1}{||\mathrm{Followers}(Y)||}$. The resulting score may be used as an indicator for X's influence and consequently for their CSC.

10.1.3 *PageRank with Topical Similarities*

A similar analysis algorithm was presented by Weng et al. in 2010 (Weng et al., 2010). They introduced TwitterRank, which takes the topical similarity between users as well as the global link structure of the network into account. The global link structure is an aspect of PageRank and allows an interpretation of a user's influence similar to the authority of a website: user A's influence is high if the sum of influences on users is high, which themselves have a high influence on others.

The topical similarity goes beyond PageRank and TunkRank, as it allows the algorithm to regard someone's influence in different topics. This is of interest for the extraction of CSC, as it looks beyond a person's general popularity (which may originate from fame) and focuses on the content of the interactions.

The TwitterRank algorithm can be divided into the three main steps: topic distillation, topic-specific relationship network construction, and topic-sensitive user influence ranking.

1 http://thenoisychannel.com/2009/01/13/a-Twitter-analog-to-pagerank/ (retrieved 2016-08-08)

1. *Topic distillation* means that for each user, a profile with topics of interests is created. This can be achieved by analyzing the content of their tweets. A direct way to infer the topic is the use of hashtags that are already in the posts. However, Weng et al. observed a low usage of hashtags in their data set (Weng et al., 2010). This is consistent with the research of Kywe et al., who analyzed 44 million tweets and found that hashtags were used in less than 8 percent of them (Kywe et al., 2012). Therefore, Weng et al. implemented Latent Dirichlet Allocation (Blei et al., 2003) topic modeling to determine the most popular topics.

   For each Twitter user a document is created that contains all their tweets. Topics are modeled as probability distributions over words. Each document is then modeled as a probability distribution over topics.

   With other measures, such as similarity measures based on the Jensen-Shannon Divergence, it is then possible to calculate the topical difference between two Twitter users. (Weng et al., 2010)

2. Based on these findings, *topic-specific relationship networks* are created under the assumption that Twitter users have different interests that are reflected in the identified topics. The resulting graph reflects the follower relationships along different topics. Twitter users are represented as vertices and edges are directed from followers to friends. (Weng et al., 2010)

3. In the last step the influence of users is measured with the *TwitterRank algorithm*. First a Twitter user is chosen at random. Then a random walk is performed along the vertices of the graph. This random walk is topic-specific insofar that the transition probability between two Twitter users is dependent on their thematic overlap (as calculated in the first step). The transition matrix is defined by (Weng et al., 2010) as follows:

$$P_t(i,j) = \frac{|\tau_i|}{\sum_{a:\ s_i \text{ follows } s_a} |\tau_a|} \cdot sim_t(i,j), \tag{35}$$

   where $P_t(i,j)$ is the transition probability regarding topic t from follower s_i to friend s_j,

   $|\tau_i|$ is the number of tweets published by s_i,

   $\sum_{a:\ s_i \text{ follows } s_a} |\tau_a|$ sums over the tweets published by all of s_i's friends, and

   $sim_t(i,j)$ is the similarity between s_i and s_j in topic t, which is defined as

$$sim_t(i,j) = 1 - |DT'_{it} - DT'_{jt}|, \tag{36}$$

   with DT' being the matrix elements of the normalized $D \times T$ matrix, where D is the number of Twitter users and T the number of topics. A matrix element DT_{it} of the unnormalized DT matrix represents the number of times a word in user s_i's tweets has been assigned to topic t.

   The transition probability captures two notions. It is dependent on the number of tweets by user s_j that are theoretically visible to user s_i because we can assume that s_j's influence on s_i increases together with her exposure to s_j's tweets. Secondly, a topical similarity between the interests of both users is taken into account by the term $sim_t(i, j)$. (Weng et al., 2010)

To address the problem of being stuck in a loop, which happens when a group of Twitter users follow only each other, Weng et al. use a transition vector E_t, with

$$E_t = DT''_{.t}. \tag{37}$$

$DT''_{.t}$ is the t-th column of matrix DT, which is the column-normalized form of matrix DT such that $||DT''_{.t}||_1 = 1$. Thanks to this extension it is possible to jump to another random vertex instead of following the graph. Now the TwitterRank $\mathbf{tr_t}$ of all Twitter users regarding topic t can be defined with the iterative equation

$$\mathbf{tr_t} = \gamma P_t \times \mathbf{tr_t} + (1 - \gamma) E_t, \tag{38}$$

with γ being a parameter between 0 and 1, which controls the probability of teleportation, and

E_t being the transition vector along which the jump is carried out.

The result is a set of topic-specific vectors that contain a person's TwitterRank. (Weng et al., 2010)

To verify their results, Weng et al. worked with a data set of the top 1000 Singapore-based Twitter users. On this data they analyzed the algorithm and identified the Twitter users with the highest TwitterRank in five categories. The fact that the number of followers of the top users varies widely is a consequence of TwitterRank's architecture that values not only the number of followers but especially the importance of the respective followers. Weng et al. also compare TwitterRank to the centrality measures In-degree (in this context the number of followers), PageRank (without topic sensitivity), and topic-sensitive PageRank. The Kendall correlation between TwitterRank and the other algorithms is 0.42 (In-degree), 0.47 (PageRank), and 0.68 (topic-sensitive PageRank). The fact that TwitterRank is most similar to topic-sensitive PageRank makes sense, because the other two measures are not directly topic sensitive. Additionally, the algorithms' performance in recommendation tasks was compared. Therefore, Weng et al. looked at randomly chosen following relationships between two Twitter users. Ten additional Twitter users without any relation to the follower were chosen. Then the link between the two original Twitter users was canceled. The task of the different algorithms was to identify the person among the other eleven who exerts the most influence on the follower. The algorithm performs well if the original friend is chosen. This is then tested for several different scenarios and different chosen follower relationships. TwitterRank achieves the best recommendation quality in most scenarios, however it does not perform better in all cases. (Weng et al., 2010)

TwitterRank and PageRank include information about the graph structure that goes beyond the pure count of performance indicators discussed by (Anger and Kittl, 2011). There is no comparison to the ratios used by (Anger and Kittl, 2011). However, the success of the algorithms, especially in the prediction task, indicate that there is relevant information in the graph structure that can be used for the prediction of a user's influence and thereby CSC. To predict concrete CSC values instead of performing a ranking, the TwitterRank or the other centrality measures of a person may be used as feature for supervised learning.

10.1.4 *Correlation of Intrinsic Metrics with a User's Influence*

In 2010 Cha et al. investigated several of the factors that we discussed in section 3.1: the number of followers, friends, tweets, retweets, and mentions (Cha et al., 2010). Their research focus lies on the correlation of these factors with user's influence, i.e. the user's "potential to lead others to engage in a certain act" (Cha et al., 2010). As influence is a constituent of contributive social capital, the investigations can be used as input for the extraction of CSC values.

The data set collected by Cha et al. is comprehensive. It contains 55.0 million accounts that are connected by 2.0 billion social links and 1.8 billion tweets. After excluding inactive accounts (who published less than 10 tweets over the time of their existence) and private accounts (whose tweets can only be seen by friends), 6.2 million users were investigated.

Some of the discussed measures fail as indicators for someone's influence, however, others corresponded to different aspects of a user's influence (Cha et al., 2010):

- A ranking based on the number of *followers* correlates with the user's potential audience. Cha et al. conclude that people with a large In-degree are popular users: an analysis of the top 20 users reveals that the accounts are primarily popular news sources, politicians, athletes, actors, and musicians. As the activity level of the followers is not considered, it does not necessarily mean that the user is influential in terms of spawning retweets and mentions. (Cha et al., 2010)

- The number of times a user is *retweeted* is seen as "the ability of that user to generate content with pass-along value" (Cha et al., 2010). An indication for this factual focus is that 92 percent of all retweets contain a URL. The top 20 users in this category are aggregation services, businesspeople, and news sites. (Cha et al., 2010)

- The number of *mentions* a person receives can be seen as the user's ability "to engage others in a conversation" (Cha et al., 2010). Less than a third of all tweets with mentions contain a URL. Therefore, Cha et al. interpret mentions as being less driven by the tweet's content and more by the author's identity. Therefore, they conclude that this measure reflects the user's name value. This is supported by the fact that most of the top 20 Twitter users in this category are celebrities (Cha et al., 2010).

- The number of *published tweets* as a sole measure for a user's influence often resulted in spammers or automated bots to be identified as most influential (Cha et al., 2010). Therefore, this measure was disregarded.

- The *outdegree*, namely the number of people a user follows (friends), suffers from the same problem and was consequently disregarded as well.

A normalization, e.g., dividing the number of retweets by the total number of published tweets, led to different results that no longer ranked the users with the highest sheer number of retweets as influential. (Cha et al., 2010)

Additionally, Cha et al. investigated the overlap between the three indicators followers, retweets, and mentions. As we have already seen from the brief discussion of the top 20 users, there may be only a small overlap. This is supported by an analysis of Spearman's rank correlation of the top ten percent of users, which confirms that the number of followers is not correlated to the other measures. Retweets and mentions correlate strongly (0.64), which indicates that people who are retweeted often are also mentioned often, and vice versa. (Cha et al., 2010)

Cha et al. also analyzed the relation between influence and the discussed topics. For that the three most popular topics in 2009 were chosen. Tweets that discussed these topics were identified with several keywords that are associated with the respective topics. A period of 60 days was investigated. Longer time periods are often problematic, as spammers will try to use the respective popular keywords. The first finding is that there is only a small overlap of people who discussed all three topics – only 2 percent of Twitter users who tweeted about one topic also mentioned the other two. The second finding is that influential users are in general not limited to only one topic but are influential over a variety of topics. (Cha et al., 2010)

Cha et al.'s findings are of high relevance for contributive social capital research on Twitter. The discussion of the metrics (the number of followers, tweets, retweets, mentions, and followees) and the corresponding types of influence is helpful for the definition of metrics and ratios as features for supervised machine learning. It is apparent that passive indicators (number of followers, retweets, mentions) by themselves are more important for the prediction of a user's influence than the metrics that can be directly altered by the user. Audience size (represented by followers), the ability to generate high quality content (retweets, mentions), and to engage in and socially maintain conversations (mentions, retweets) are properties expected of people with high CSC and value for their network. It was not investigated if a combination of these features would further improve the results. This was done by (Hadgu and Jäschke, 2014) (see next section). The finding that a normalization of the features does not directly correlate with influence is opposed to the findings by Anger et al. (Anger and Kittl, 2011), who were able to correlate normalized features to a user's influence. For CSC research this means that further research would be helpful, ideally with the help of CSC ground truth values.

10.1.5 *Expertise on Twitter*

A method that deals with the expertise of a person was presented in Hadgu and Jäschke's 2014 paper (Hadgu and Jäschke, 2014). They investigate how various algorithms can be leveraged to identify researchers, who can be regarded as experts in their fields, on Twitter. Hadgu and Jäschke used support vector machines (like Rao et al.), random forests, classification and regression trees, and logistic regression for the classification. The task of identifying researchers is one of binary classification, which is different from inferring CSC scores for each user. Therefore, we will discuss the parts of their research that are of relevance for both tasks — the used data set, the implemented methods and set the results in relation to CSC prediction.

For the analysis, two data sets were collected and then combined. The first one consists of a random crawl of 1 million Twitter users. As the percentage of scientists among Twitter users can be assumed to be comparable to the average population, most of these users are not likely to be researchers. The second data set was selected with the intention of sampling a group of scientists. To achieve this task, the followers of the accounts of scientific conferences (e.g., @www09) were crawled. (Hadgu and Jäschke, 2014)

To create a set with ground truth data (here: scientist yes/no), the users in the second set were matched against authors from DBLP, a computer science bibliography hosted at the University of Trier. The resulting list contains the positive examples. To create a list with negative examples (no scientists) the first data crawl was used. All users that follow any of the conferences and their followers were subtracted because of the increased likelihood of them being scientists. From the resulting list a random subset was drawn and used as ground truth set for people who are not scientists. (Hadgu and Jäschke, 2014)

In the next step, Hadgu and Jäschke defined features for the machine learning algorithms. Some features were obtained from profile information, like name, location, URL (e.g., if the top level domain is ".edu"), description (e.g., if there are specific keywords like "phd"), and total number of tweets, followers, and friends. Other features were selected from the content of the tweets, e.g., the number of tweets and retweets, different ratios (retweets to tweets, tweets containing URLs/hashtags), as well as distinct hashtags (e.g., conference names). The distinct hashtags were selected from the ground truth data set that only contains scientists. 1,872 hashtags were selected in that way. (Hadgu and Jäschke, 2014)

With these features, various classification algorithms can be compared on the data set. Hadgu and Jäschke performed stratified 10-fold cross-validation to train the models on 2,000 of the positive and all the negative profiles. The performance of all algorithms (support vector machines, random forest, classification and regression trees, and logistic regression) was satisfying. Random Forest performed best along all performance measures (precision: 0.96, recall: 0.92, F1-measure: 0.94, accuracy: 0.95, and true negative rate: 0.97). The other algorithms performed between 0.88 and 0.90 (precision), 0.87 and 0.90 (recall), 0.88 and 0.90 (F1-measure), 0.89 and 0.91 (accuracy), and 0.90 and 0.92 (TNR) and are, therefore, all useful for the analysis. (Hadgu and Jäschke, 2014)

The importance of the different features was measured by their mean decrease accuracy, which reflects how much using the feature in the classifier reduces the classification error. The most important feature was the number of tweets, followed by the number of tweets with scientific hashtags, the friend/follower ratio, and whether the user description contains keywords. (Hadgu and Jäschke, 2014)

We are looking for methods to infer CSC values for each user. For this task we can leverage several findings from Hadgu and Jäschke's research. For the CSC prediction, the machine learning classifiers need to be replaced with regressors and the ground truth values need to be replaced with CSC values. The used data set and the selected features can be used in the same way. The findings about the importance of the features can likely be transferred to the regression task, as the both tasks are similar in nature and use the expertise of a person as input. Consequently, the number of tweets, the number of tweets with scientific hashtags, and the friend/follower ratio should be used for CSC identification. Random forest, classification and regression trees, support vector machines, and logistic regression may all be used for regression analysis because they were successful for the classification.

10.1.6 *Summary: CSC on Microblogging Platforms*

There is a variety of papers that deal with analyzing and categorizing users on the microblogging network Twitter. Some of the discussed procedures can be used directly for the extraction of contributive social capital, or its constituents influence and expertise. If a ground truth labeling with CSC values were available, one could use regression methods for the extraction of CSC from the presented data sets, as well as a mapping of the expertise and influence scores on the CSC per person. This would allow the creation of a structured CSC value that attributes parts of the CSC to a user's influence and parts to the user's expertise. In the following, the key takeaways of the discussed investigations are summarized:

- There are several useful features (number of tweets, number of followers, etc.) and ratios thereof (follower/following ratio, retweet and mention ratio, etc.) that can be either directly used to identify influential people, or as an input for analysis algorithms. Anger et al. demonstrated, e.g., how these ratios relate to a person's influence (Anger and Kittl, 2011), which is an attribute of social capital. A variety of additional features can be used as input for classification and regression algorithms (e.g., different linguistic features, like the use of parts of the user-profile by Hadgu and Jäschke).

- There are two main ways to crawl the data for the analysis: random crawls and specific crawls of sub-networks to create data sets of people with specific properties (e.g., Hadgu and Jäschke crawled scientists by selecting the followers of the accounts of scientific conferences to investigate expertise on Twitter). Both methods present advantages for CSC research. Random crawls extract a subset of the full network that shares the same characteristics and can, therefore, be used for the analysis. Specific crawls can be used as addition to create labeled ground truth data sets.

- Ground truth labels on the data sets can either be obtained by manual labeling (e.g., to identify a person's regional origin based on language differences), or by selecting the data from Twitter lists (e.g., participant groups like the NRA, see Rao et al.). Additionally, one can match the user names with external databases (e.g., with scientific databases, see Hadgu and Jäschke).

- Various algorithms can be used for the analysis: Intrinsic features (like the number of followers in Cha et al.), ratios (e.g., the Retweet and Mention Ratio by Anger et al.), PageRank-like algorithms (TwitterRank by Weng et al.), and supervised learning algorithms (support vector machines, decision trees and random forests, e.g., by Hadgu and Jäschke). All procedures may be used to identify

CSC or its constituents on different microblogging data sets. Single features by themselves were less successful than methods that take more information about the network into account (see (Weng et al., 2010)'s comparison of PageRank and TwitterRank vs. the count of followers).

These findings were used as input for the experiment that investigated CSC extraction from Twitter (section 12.6).

10.2 SOCIAL NETWORKING PLATFORMS

Social networking platforms are another data source that is relevant for the analysis of social capital.

Analogous to the procedure in the last section, we discuss previous work that deals with the extraction of people's attributes, or the classification of people along different properties that are either directly relevant for social capital or indirectly, in the form of methods that can be adapted, e.g., with new ground truth data.

10.2.1 *Measuring Influence with Direct Network Features*

In 2013, S. Hassan published a paper on identifying criteria for measuring influence of social media (Hassan, 2013). The main purpose of the study was to identify practical measures for influence in different social media communities. He investigated the intrinsic features as discussed in section 3.2 and defined seven features as most relevant. Those were then categorized along three different dimensions of influence:

- Recognition
  - Number of likes
  - Number of subscribers/followers/friends
- Activity generation
  - Number of posts
  - Number of received comments on written posts
  - Number of shares of the user's posts by others
  - Number of in-links (number of times a user or her posts are referenced)
- Novelty
  - Number of outlinks (how often resources outside of the website are linked)

These findings were based on a content analysis of social media. However, no quantitative assessment was used to verify the findings, e.g., by comparing against ground truth data. The identified criteria are plausible and are in line with the previously discussed analysis on Twitter (compare for example Cha et al. (Cha et al., 2010)). This indicates a usability for social capital extraction of the named features, e.g., as input for supervised learning algorithms.

10.2.2 *Formularizing Social Influence*

J. Bentwood introduced the "*Social Media Index*." (Bentwood, 2008) This index uses input from different types of social media to estimate the overall online influence of a person. The input is based on the six following criteria (Bentwood, 2008):

- Number of friends on Facebook
- Google rank, in-bound links, subscribers, Alexa rank, content focus, frequency, number of comments of the person's blog
- Number of friends, followers, and updates on Twitter
- Number of contacts in LinkedIn
- Number of photos by and with the user on Flickr
- Favorites on Digg and del.icio.us

This multi-level approach is comparable to the Klout score, which is discussed in the next section.

Based on these findings, Bentwood at al. describe a person's influence/online presence in the form of an expression (Bentwood, 2008):

$$\frac{\text{Volume and Quality of Attention} \times \text{Time}}{\text{Size and Quality of Audience}} \tag{39}$$

The components of the expression are not clearly defined, which means there is no formal model of how to directly assess the quality of attention, for example. It can, however, be approximated with intrinsic features and tailored to the goal of the analysis.

For the extraction of contributive social capital, the following metrics may be used as input for the expression. The volume and quality of attention could be measured with the number of friends a user has and the number of likes a user receives — both metrics that were described by (Hassan, 2013) as measures for recognition. An interesting part of the equation is the normalization in the denominator because it allows to take the audience's properties into account, which is related to the PageRank based approach by (Weng et al., 2010). By doing so, one can value the input of people proportionally to their respective contributive social capital. The size of the audience can be determined either by assessing the number of views of specific posts, or other metrics like the overall number of posts (to normalize for posts without reaction), the number of friends/followers (which reflects the ratio of responses), or the size of the network (which reflects the ability to reach people).

10.2.3 *Commercialized Influence Calculation with Supervised Learning*

Rao et al. described an approach to measure a person's online influence over several platforms (Rao et al., 2015). They trained a supervised learning algorithm to calculate the Klout score of a person. This score represents their influence on a logarithmic scale from 1 to 100. The used data sources are social networking platforms, like Facebook, Google+, LinkedIn, microblogging (Twitter and Instagram), and other web portals, like Wikipedia. Most of the other publications discussed in this paper use static data sets, which means the data was crawled once and then analyzed. Contrary to that the Klout score analyzes up to 45 billion interactions of up to 750 million users on a daily basis. In the following, we briefly describe the method, feature types, and validation mechanisms to make this commercial software comparable to other approaches. The algorithm works along several steps:

- A feature vector is created for every user for each network or community that is used as data source.
- Each feature has a different weight for the computation of the overall influence score of the user.
- The feature weights were determined with a supervised learning model trained on labeled data. For the calculation, Rao et al. used non-negative least squares regression. The weights are not published.
- The result from the different data sources is hierarchically combined to create one score per user.
- In the last step the resulting value is scaled to a value between 1 and 100 on a logarithmic scale.

Users have to sign up to use the system. Once someone is registered, the required data can be extracted with different APIs based on the granted permissions. The Twitter data is extracted from Mention Stream[2]. To obtain the ground truth data for the supervised learning algorithm, Rao et al. presented evaluators with pairs of people. The evaluators knew the people from their network. They labeled the person they found to be more influential. In order to prevent mistakes, all pairs were evaluated by several people. In total, over 1 million evaluations were used to determine the weights.

A total of 3,600 features can be used by the algorithm. Rao et al. do not publish a complete list; however, the following features are included:

- Analysis of the interaction graph between users (graph structure and reaction types)
- Profile information (e.g., job title from LinkedIn)
- Number of friends on Facebook
- PageRank derived from Wikipedia
- Number of news articles that mention the user

Rao et al. used two different approaches to validate their findings. Both methods can be leveraged to design verification mechanisms for CSC extraction.

For the first experiment, which was conducted over the course of a year, various users were encouraged to write about certain perks they received. The reaction of

2 https://gnip.com/sources/Twitter/ (retrieved 2016-08-29)

their audience was measured. Plotting the average number of reactions to a post over the Klout score of the respective user, produced a monotonically increasing curve. This is a strong indication that there is a correlation between the Klout score and a user's ability to spread information.

Another approach for validation is to study the overlap with other lists of influential people. In (Rao et al., 2015), Rao et al. publish two such comparisons between the Klout score and the ATP Tennis Player Ranking and the Forbes' list of most powerful women. The reasoning behind choosing both lists is not explained in (Rao et al., 2015). The overlap is determined with the normalized discounted cumulative gain metric (nDCG, see equation 40, 41) and the calculated scores of 0.878 and 0.874 indicate a large overlap between both lists and the Klout ranking.

$$nDCG_p = \frac{DCG_p}{IDCG_p}, \tag{40}$$

where p is a particular rank position and $IDCG_p$ is the ideal DCG through position p.

$$DCG_p = \sum_{i=1}^{p} \frac{2^{rel_i} - 1}{\log_2(i+1)}, \tag{41}$$

where rel_i is the graded relevance at position i.

The Klout score and the CSC system both aim to visualize an individual metric for each participant.

For the task of contributive social capital retrieval, one can consequently leverage several of these findings. Defining features (compare (Hassan, 2013) and (Bentwood, 2008)) that are then used in supervised learning is a method that can be used to identify a person's contributive social capital with the help of different ground truths. Similarly, the data sources and the validation mechanisms can be adapted for CSC extraction.

10.2.4 *Network Analysis with Centrality Measures*

Applying any of the centrality measures that were presented by (Sun, 2011) and in section 4.1.2 to a graph, extracted from a social networking platform like Facebook, yields aspects of the importance of a user within the investigated graph structure. As seen on Twitter, there is a correlation between a user's centrality measure and their influence (Weng et al., 2010). Depending on how the extracted graph was created, this allows additional conclusions about the user's CSC. It is, e.g., possible to refine this methods in order to obtain a content- or domain-dependent score, which is achieved by reducing the network to edges between nodes that deal with certain topics (e.g., by only regarding discussions on computer science) and repeating the centrality analysis process. Additionally, the centrality measures can be used as features in supervised learning algorithms that predict CSC.

10.2.5 *Trust in Online Social Networks*

Another aspect of CSC was discussed by Ziegler at al. in (Ziegler, 2009). The focus of this work is the assessment of trust between agents from online social networks. For their investigations, they distinguish between local and global trust. Local trust is the personal feeling between different agents and the degree to which they trust each other. Global trust is the reputation of a person that can be inferred from information in the network.

Most of the algorithms used for the extraction of global trust work similarly to PageRank; the global trust of an agent is related to the global trust of the agents linking to him. Examples for such metrics are EigenTrust by Kamvar et al. (Kamvar et al., 2003) and PowerTrust (Zhou and Hwang, 2007). Local trust metrics, on the other hand, "take the agent for whom to compute trust as an additional input parameter and are able to operate on partial trust graph information." (Ziegler, 2009) This allows for an individual trust assessment. In order to obtain a usable score, local trust metrics need to leverage structural information defined by personalized webs of trust. (Ziegler, 2009)

For CSC analysis, we are interested in an objective measure for the social capital an individual adds to the network and not primarily in the subjective assessment of parts of the network. Consequently, global trust metrics are of higher relevance. They are similar to PageRank, an algorithm whose importance for contributive social capital research was already underlined (compare section about TwitterRank (Weng et al., 2010)). The resulting lists may be used as features for CSC prediction with supervised learning.

10.2.6 *Summary: CSC on Social Networking Platforms*

As described in this section, there is various approaches that can be used in order to investigate social capital on social networking platforms. The discussed methods can either be used directly (if the investigated properties are constituents of CSC) or indirectly (by adapting the procedure, e.g., with the help of a new ground truth). The key takeaways of this section are summarized in the following:

- Similar to microblogging, there are a number of intrinsic performance indicators that can be used as a measure for the extraction of CSC or its constituents. An overview of different metrics and their categorization was offered by S. Hassan (Hassan, 2013). This is similar to the extraction methods used by Cha et al. for the microblogging service Twitter (Cha et al., 2010).

- In order to tailor the extraction method to the specific interests of the study (e.g., social capital), these measures can be combined to create a better depiction of reality. An exemplary equation was provided by J. Bentwood in (Bentwood, 2008). This makes it possible to create ratios similar to the ones discussed by Anger et al. in the previous section (also see (Anger and Kittl, 2011)). We offered an interpretation of how this equation might look for CSC, based on the previous work by (Hassan, 2013).

- The Klout score by Rao et al. (see (Rao et al., 2015)) is an example for a machine learning algorithm with several different data sources. This measure uses supervised learning and was tested successfully in several experiments. It is related to the CSC system in the way that it aims to assess an individual score for each participant.
- We also reviewed previous work about degree, closeness, betweenness, eigenvector, and PageRank centrality and ways how they can be leveraged for CSC analysis. The resulting values can also be used as input for supervised learning algorithms.

These findings were used for our experiments about the extraction of CSC from Facebook (section 12.5) to adjust our analysis to the platform.

10.3 DIRECT COMMUNICATION

Whereas Twitter is limited to 280 characters per tweet and posts on social networking platforms are usually also kept short, there is no relevant limit for email communication. Additionally, emails are often used for the exchange of information. This makes the analysis of email networks so compelling for the extraction of CSC and its constituent expertise. The research already conducted in this field can be broadly categorized along two main targets: expert finding and social status/hierarchy extraction. Both topics have an overlap with social capital. Expert finding is connected to assessing the expertise of a person, which is a constituent of contributive social capital (see section 2.2). While the direct correlation between social capital and social status remains to be investigated, we still want to discuss it in this section because the social status of an individual can be regarded as consequence of this person's CSC. In the following, we discuss several exemplary articles to present the ongoing research by describing the method used, the data set and verification, as well as the findings in relation to CSC.

10.3.1 *Expert Identification Using Search Algorithms*

Zhang et al. compared algorithms designed for expert identification (Zhang et al., 2005). They evaluated searching strategies from three different families: general computational (breadth first search, random walk), network structure based (best connected, weak tie, strong tie, cosine similarity, hamming distance), and similarity based (information scent). These algorithms are used to create a match between a person with an information need and the best expert to provide the required information. The matching process itself is of little direct interest for social capital analysis. However, the discussion of the experimental verification method is relevant, as it may be leveraged in social capital extraction research.

The investigations were carried out with the help of the Enron data set, which contains about 500,000 emails from 147 employees. For the analysis a graph was constructed with 32,766 messages as edges between the 147 employee-nodes. An information profile for each user was created by using TF-IDF for all messages sent and received by the user. It is an approximation that these represent the person's information space but there was no feasible way to obtain better profiles of the former Enron employees.

For the evaluation Zhang et al. defined different queries and assigned them randomly to a node in the network. The eight search strategies were used to find a match between the 'information seeker' and the person who was most likely able to provide the answer. The match was evaluated using a standard TF-IDF measure, therefore, the person with the exact same combination of keywords within his or her information space would receive the highest score. The highest success rate was achieved with breadth first search, closely followed by information scent, best connected, cosine similarity, and hamming distance.

Even though the focus of this work lies on the search algorithms, there are several aspects that can be leveraged for CSC research. By counting the number of times a person was identified as expert in a search, one could create a metric for each person's topic-sensitive expertise. As this metric would respond to how often the person helped others in the network, it should be correlated to the person's CSC. This metric could also be used as input to supervised learning algorithms. Additionally, the data set is publicly available and can be investigated further.

10.3.2 *Expert Identification with Bayes' Theorem*

Another article that deals with the topic of expert identification based on email data was published by Balog et al. (Balog, 2006). They use Bayes' theorem to identify a person with the highest probability to be an expert according to a defined query. This allows to identify the top experts in different topics — which could be used as input for CSC identification.

The probability that a candidate ca is an expert in topic q can be calculated with the help of Bayes' theorem:

$$p(ca|q) = \frac{p(q|ca)p(ca)}{p(q)} \tag{42}$$

Balog et al. used the resulting probability $p(ca|q)$ to rank the experts for each topic q. For that, $p(ca|q)$ was calculated with two different models. Model 1 computes $p(ca|q)$ by constructing a candidate model from all documents associated with candidate ca by collecting all term information from them. This candidate model is then used to represent the candidate in the query. The approach of model 2 is different. It assumes that q and ca are conditionally independent and that their relation can be described through document-candidate associations. With the help of these document-candidate associations the person is determined who is most strongly associated with the documents that describe the topic. This requires investigating each document (e.g., email) to identify who is associated with it.

In both models, each candidate is scored by aggregating over all email documents associated with the candidate. The topical overlap between the emails and the query is calculated with language modeling.

For experimental evaluation Balog et al. used the email W3C data[3], which comes with a list of candidate experts, topics, and relevance assessments for these topics. On

3 http://research.microsoft.com/en-us/um/people/nickcr/w3c-summary.html (retrieved: 2016-09-15)

this data, the association of a person to the email was investigated. The findings are summarized in the following.

Model 2 generally outperforms model 1. Generally, there are four ways in which a person can be involved in an email: as sender, recipient, in cc, or within the message content. Balog et al. found that a person who is the sender of an email has the highest correlation to that person actually being the expert on the content of the email. (Balog, 2006) Consequently, this feature should be included in future CSC investigation with the help of regression from features.

For CSC extraction the correlation of $p(ca|q)$ with CSC ground truth values can be investigated to assess to what extent $p(ca|q)$ can be used as input or proxy for the contributive social capital of a person.

Expert Identification with Communication Patterns and Message Content

Campbell et al. researched ways to identify experts in email communication data sets (Campbell et al., 2003). Their objective was to identify experts using communication patterns as well as message content.

The algorithm presented by Campbell et al. can be divided in three different steps. At first, all emails related to a topic are collected. This is achieved by keyword retrieval and clustering. In the second step a communication graph is created in which the nodes correspond to people and the edges are directed from email sender to recipient. In the third step a modified version of the HITS graph-based ranking algorithm is implemented.

With the reasoning that people send their questions to others who, in their opinion, are able to provide an answer, Campbell et al. use the authority score as a measure for someone's expertise.

In order to evaluate this approach, Campbell et al. compare it to a query term frequency approach on two email data sets. Both algorithms were compared to a ground truth expertise score between 1 and 10. The query term frequency approach is similar to the one discussed by Balog et al. as it ranks a person according to the number of emails sent on a specific topic. The person ranked the highest in a specific topic with this algorithm consequently sent the most emails related to this topic. One data set was extracted from a research organization (15 participants and 13,417 messages) and one from a software development organization (9 participants and 15,928 messages). The expertise score was created by averaging over ratings from other candidates. People who received an average expertise score of over 7 were marked as experts. (Campbell et al., 2003)

The comparison was conducted with the help of several metrics: percent of correctly identified experts, percent of false alarms, precision, and detection (which combines correct identification, incorrect identification, incorrect rejection, and correct rejection). Overall the HITS algorithm did perform significantly better than the query-term frequency approach. The detection score of HITS was 0.39 (0.63) vs. 0.28 (0.44) for the research organization (software development organization). A similar advantage was observed for precision with 0.52 (0.67) for HITS vs. 0.38 (0.50) for the content-based approach. The percentage of false alarms was with 0.35 (0.55) vs. 0.71 (0.64) — again significantly better for HITS. Only the percentage of correctly identified experts was better for the content-based approach, with 0.38 (0.33) vs. 0.44 (0.31) (Campbell et al., 2003).

These findings are relevant for CSC extraction from email communication, as they underline that there is information about a person's CSC that can be extracted via analysis of the graph structure with the HITS algorithm in a way that is superior to straightforward term-frequency approaches. This also supports the findings by Weng et al. (Weng et al., 2010) on Twitter, who successfully identified influential users with other centrality measures. The authority scores calculated with the HITS algorithm can consequently be used as input for CSC retrieval with supervised learning.

10.3.3 *Social Hierarchy in Email Networks*

Social hierarchies were investigated by Rowe et al. (Rowe et al., 2007). They presented an algorithm to automatically extract social hierarchy structures from electronic communication networks.

Like Zhang et al. they worked with the Enron email data set, which contains communications from employees along several hierarchies. The employees can be cross-referenced against the organizational charts.

The algorithm works as follows. For each user several features are analyzed to create a ranked user list. The first analyzed aspect is the information flow between users. The investigated features are:

- number of sent and received emails (hypothesis: more important people need to communicate more),
- average response time (hypothesis: the faster user i answers to user j's request, the more important j is to i.)

The next aspect Rowe et al. investigated is the communication network graph that was created in a way that all users are vertices that are connected by undirected edges if they exchanged at least N emails. In this context the following features were calculated for each user:

- Number of cliques an account is associated with,
- raw clique score that also takes the size of the cliques into account,
- weighted clique score that regards the importance of other clique members (determined by the number of emails and average response time),
- centrality measures: degree-, closeness-, and betweenness-centrality as well as the clustering coefficient, and the
- hubs-and-authorities importance of each user.

These individual metrics are then weighted and a social score between 1 and 100 is computed along which the users are ranked. (Rowe et al., 2007)

The results were compared with the organizational chart and overlap significantly. Rowe et al. were able to predict the social hierarchy of the 54 traders in the North

American West Power Traders division of Enron with great accuracy. A higher standing within an economic setting usually translates to better access to resources. Therefore, these findings can be used for the information access (expertise) aspect of CSC and the correlation of the social score and ground truth CSC values investigated.

10.3.4 *Social Status in Email Communication*

Bird et al. investigated the extraction of social status of email users (Bird et al., 2006). As data set they used the email list of an open source software (OSS) development project. Over the course of seven years Bird et al. collected roughly 100,000 messages of more than 1000 users.

To evaluate a person's importance within the network Bird et al. used the centrality measures betweenness, in-degree, and out-degree. According to these measures, the developers among the OSS group are higher in status than the non-developers, which is reasonable regarding the nature of the data set. Within the group of developers there is also a significant correlation between the centrality measures and the number of committed changes to the code. This indicates that the developer's importance correlates to these measures. This indicates that the centrality measures are relevant for CSC extraction from email communication. This underlines the findings by (Campbell et al., 2003) and (Weng et al., 2010) that information about the constituents of CSC of users is stored in the graph structure of communication networks, which can be accessed via centrality measures.

10.3.5 *Summary: CSC in Email Communication Networks*

The key findings of this section can be summarized as follows.

- There are less features available than on microblogging or social networking platforms, as, e.g., there is nothing similar to the "likes" on Facebook. Features that are relevant for the CSC of users are the number of emails sent and received, response time, clique measures, and centrality measures (as presented by Rowe et al. (Rowe et al., 2007) and Bird et al. (Bird et al., 2006)). This is similar to the previous findings in other data sources.

- The success of centrality measures (e.g., betweenness centrality, HITS) for the prediction of CSC constituents indicates that there may be relevant information within the graph structure of the investigated networks (Campbell et al., 2003), (Rowe et al., 2007), (Bird et al., 2006)

- For expert finding, an established approach is to compute an overlap between a defined query and the information space of a user. This can be done with TF-IDF (Zhang et al., 2005) or with other language models (Balog, 2006).

- The standard data sets used for the investigation of email communication stems from companies or research organizations, e.g., the Enron email corpus.

These findings are of relevance for the expertise aspect of contributive social capital, as the content-focus of email communication make it ideal for expert identification as well as social status and hierarchy assessment. The presented approaches and algorithms may be used directly for the extraction of the CSC constituent expertise, or indirectly by adapting them for CSC extraction (e.g., using the discussed features and centralities as input for supervised learning algorithms with CSC as ground truth values). There are many similarities between email communication and other means of direct communication, e.g., similar features (sender, recipient, length of the message, message content, etc.) and graph structure. Due to these similarities, it is likely that the findings can be extended to cover these other networks (e.g., WhatsApp communication), as well. When doing so, it is important to adjust for the differences, e.g., by adding new features (groups in WhatsApp) or removing them (CC recipients in emails is not available in all means of direct communication).

We did not investigate CSC extraction from direct communication in our experiments because we assume that people do not want their private messages analyzed within the CSC system. The observations in this section, nevertheless, support the findings on the other data sources and facilitate, if necessary, research about CSC extraction from direct communication.

10.4 CITATION NETWORKS

As discussed in section 2.1.2, a key aspect of someone's CSC is their knowledge and expertise. By providing an expert's opinion on various topics, people can help their surrounding network making decisions and building opinions. This increases their value within the network. A type of networks that directly deal with expertise are scientific citation networks. In section 3.4, the nature of citation networks were briefly presented and citations were mentioned as direct performance indicators. However, there is a number of methods that go beyond the pure count of citations. In the following, publications about citation network analysis are discussed and the findings are set in relation to contributive social capital extraction.

10.4.1 *Hirsch Index, g Index, and i10 Index*

There is a variety of indices that are used throughout the scientific world to measure a scientist's importance.

A popular index was introduced by J. Hirsch and sets the number of publications of an author in relation to the number of times they are cited: "the index h is defined as the number of papers with citation number >h, as a useful index to characterize the scientific output of a researcher" (Hirsch, 2005). This makes the resulting score robust against outliers, e.g., scientists who publish many papers with little positive feedback (e.g., the count of citations), or scientists who publish only one paper which is cited often but do not publish anything else. J.Hirsch verified the validity of the h-index on a data set of physicists and found that scientists with a high h-score are indeed prominent. He concludes that h provides an estimate of their "importance, significance, and broad impact" (Hirsch, 2005).

Abbasi et al. investigated the Spearman correlation between an author's social capital and their h-index and citation count (Abbasi et al., 2014). They found a positive correlation of 0.57 (h-index) and 0.44 (citation count). This indicates that the h-index is a better measure to approximate someone's social capital than the pure count of citations. (Abbasi et al., 2014)

With the g-index, L. Egghe presented a variation of the h-index in (Egghe, 2006). It is a refinement of the h-index insofar that it also regards the number of publications of an author and the respective citations. Additionally, it acknowledges very successful articles, which is important to identify the most successful experts. The value of g is defined in a way that the top g articles of a scientist were cited (altogether) at least g^2 times.

Google Scholar introduced the i10-index, a straightforward measure that reflects the number of articles with at least ten citations[4]. By ignoring the publications that promoted few responses, this metric focuses on popular articles. However, this also means that lots of publications that receive only low levels of attention are disregarded. The same happens with new publications that are still rising in popularity but have not passed the threshold yet.

The indices could also be implemented for CSC investigations on collaborative Q&A portals or microblogging platforms.

These indices can be used as a measure for the expertise aspect of contributive social capital, optionally after categorizing along different topics. They can also be used as features for supervised learning algorithms.

10.4.2 *Trust and Reputation with PageRank and Graph Analysis*

Jøsang et al. recognized that the standard metric "count of citations" is an indicator of quality and reputation, however, they also see shortcomings, like the fact that only positive referrals are possible (Sang et al., 2007). In order to counter this issue, they suggest the investigation of scientific citation networks with a PageRank like algorithm. Others later realized this, among them Yang et al. (Yang et al., 2010).

In 2010 Yang et al. implemented several analysis methods for citation networks, among them PageRank (Yang et al., 2010). The goal of their studies is closely related to identifying the CSC of a scientist. Based on scientific network data, Yang et al. propose a method to estimate a researcher's importance, contribution, and reputation as well as a ranking according to these findings. We discuss the methods they investigated, the data sources used for evaluation, and their results.

The focus is the investigation of three methods that implement properties of both, topical link analysis as well as citation graph analysis:

- Multi-type citation graph analysis (including relationships of authors, papers, affiliations and publishing venues) combined with content-based analysis,
- heterogeneous PageRank model, and
- topical PageRank model.

For the multi-type citation network analysis, graphs were created with the nodes being either authors, papers, affiliations, venues, or combinations thereof. The (di-

4 http://googlescholar.blogspot.de/2011/11/google-scholar-citations-open-to-all.html (retrieved 2016-07-04)

rected) edges between those social actors are all types of possible relations, like co-authorships, citations, publications in venues.

Heterogeneous PageRank could be directly applied to these multi-type citation networks by evenly distributing a node's authority among the connected nodes (authors, papers, affiliations, or venues). However, Yang et al. reason that this approach does not correctly represent the directional probabilities between the different types of social actors. Therefore, they propose the following approach, which assigns different propagation probabilities to the different types of out links (Yang et al., 2010):

$$PR(i) = (1-d) \sum_{j:j \to i} \beta_{ji} \frac{PR(j)}{O(j)_{type(i)}} + d\frac{1}{N}, \text{ where} \tag{43}$$

- j is the node pointing to node i.
- d is a random jump.
- β_{ji} represents the propagation probability from node j to i and is always identical for jumps to nodes of the same type. Also $\sum_{type(i)} \beta_{ji} = 1$ holds true, which means that the probability to jump to nodes of one of the different types is always normalized to 1.
- $O(j)_{type(i)}$ is the number of outlinks from node j to nodes of the same type as i.
- N is the total number of notes of the network.

The next step in their analysis was the introduction of a topical element by introducing different probabilities for the random jump. This allows following a link within the network, jumps to a random link within the same topic, or random jumps to another node in a random topic. This takes the notion into account that an author's expertise in one subject is not directly transferable to other topics.

For experimental verification, Yang et al. used a crawl of the ACM digital library[5] which consists of 172,891 distinct web pages that represent different publications (including publishing venue, authors, affiliation of each author, and citation references). This data was then used to create the aforementioned graphs.

The ground truth data, i.e. rankings of authors in different topics, was needed for experimental comparison of the different methods. This data was obtained in three ways.

- Authors were labeled with the help of the Libra rankings of conferences they attended.
- The ACM database labels were used (ACM fellow, ACM distinguished, ACM expert).
- Four judges reviewed the Google Scholar and other search results and labeled the scientists based on their assessments.

5 http://dl.acm.org/dl.cfm (retrieved 2016-09-09)

Normalized discounted cumulative gain (nDCG, see equations 40, 41) was then used as a metric to compare the three different methods for queries in 23 categories. The main results can be summarized as follows. The multi-type citation graph (with nodes being authors, papers, venues) led to better results than graphs that only use one of the types as nodes. Therefore, the affiliation between authors and the venues, where the papers are published, also provide useful information. This finding can be expanded to the other data sources as well. On social networking or microblogging platforms, e.g., a multi-type graph could include the user's posts or tweets as nodes, additionally to the users themselves. Using the topical PageRank improved the performance. (Yang et al., 2010)

These methods can be leveraged to extract the knowledge aspect of social capital of scientists. The determined rank can be used as input for CSC analysis.

10.4.3 *Expert Identification in Communities*

In 2012, Su et al. published an article about identifying the top k experts within a community (Su et al., 2012). They put special focus on the topic context in order to capture all aspects of an author's expertise. Additionally, it was investigated how these search results could be diversified. This task by itself is of little relevance to CSC retrieval, but the methods, the data set, and the results can be leveraged to assess the knowledge aspect of social capital in social networks.

The first step in Su et al.'s method was to identify topic information. They discussed several methods, from using pre-defined categories, via user assigned tags, to statistical topic modeling. The chosen method is the author-conference-topic (ACT) topic model that was presented by Tang et al. in 2008 (Tang et al., 2008). This model allows to calculate the probability of a connection between an author or paper to certain topics.

Additionally, Su et al. defined an objective function that evaluates the precision with which the algorithm ranks the experts vs. ground truth data. The performance was influenced by the weights that were assigned to the features which are publications, h-index of the author, and the language model-based relevance score. The weights were determined by maximizing the objective function with the help of a standard greedy optimization algorithm. After the weights were set, the algorithm could be used to determine the top k experts for each query.

The data set used for evaluation was an unnamed scientific network with more than 1 million authors and 2 million papers. In order to create the ground truth subset, a similar approach to the one presented by Yang et al. was used. For each investigated search query the best conferences were chosen. Then the chairs and committee members of these conferences were ranked according to appearance and h-index. The resulting top 100 people in the list were labeled "experts" and chosen as ground truth.

Su et al. compared their algorithm against several alternative methods (language modeling, topic modeling, and random walk) and found that it improved the ranking performance.

For CSC retrieval both contributions can be leveraged; the topic-model to determine the topical context of a user's expertise, and the ranking using an objective function and an effective algorithm to solve it. The resulting ranks may be used as input for supervised learning to determine a user's contributive social capital.

10.4.4 *Citation Network Analysis with Mendeley Metrics*

In 2013, Li et al. investigated the social media platform Mendeley[6] that allows users to archive, comment, and share their paper bibliographies (Li and Gillet, 2013). This yields several new features that can be used for the analysis, for example the total number of readers of an article. The focus of Li et al. was to identify influential scholars along two dimensions: academic and social impact. In this section the data set, methodology, and findings are briefly presented.

The data set was extracted from Mendeley in 2012 and comprises approximately 1 million user profiles, 100,000 papers, and various types of connections between the users, like co-authorship, contacts, or group memberships.

The academic impact of a scholar is high if his or her publications are popular and received well by peers. The assumption of Li et al. is that the most relevant criteria for this is the number of times a paper is read. Considering that many people read papers and do not necessarily cite them, this is a reasonable extension to the previously discussed features. Based on the number of readers they defined three different metrics (Li and Gillet, 2013):

- Total number of readers (indicates the author's overall influence),
- maximum number of readers per paper (to identify scholars with few but very visible papers), and
- R-index (an R-index of n represents that the author has published n papers with at least n readers. This is a measure of a scholar's productivity and impact and similar to the h-index for citations.)

Li et al. plotted these three metrics over the respective number of scholars and identified power law patterns. This indicates that there are only a few scientists that can be regarded as real influences, whereas the majority have low scores. Another observation was that all three measures yield different scientists in the top one, respectively top ten percentile[7]. Therefore, the three measures can be used to describe different types of influential scholars. (Li and Gillet, 2013)

For the analysis of the social influence of an author Li et al. were interested in a user's connections within the network and their ability to control information flow. They used the three centrality measures that we discussed in the context of social networking platforms: degree-, closeness-, and betweenness-centrality. Li et al. used scientists with at least one published paper as nodes and all interrelations between them (e.g., co-authorships) as edges. The results were several sub-graphs, as there is little to no overlap between different scientific fields. The centrality measures were then used to investigate local research communities. These were created using a Latent Dirichlet Allocation (Blei et al., 2003) topic model to generate topic vectors for

6 https://www.mendeley.com/ (retrieved 2016-09-02)

7 The overlap values for the top 10% (top 1%) of scientists for the different measures are: R-index vs max reader count: 24.4% (7.9%), R-index vs total reader count: 45.5% (29.8%), max reader count vs total reader count 77.7% (62.5%).

each paper and cosine similarity to compute the similarity between the topic vectors of the papers and of 25 characteristic search queries. Afterwards, for each node a single influence value was created by calculating the Euclidian norm of the normalized scores of the three measures. The scholars identified in this way are mainly young researchers — opposed to those identified to have high academic influence, who tend to be older. Li et al. also used Spearman's rank correlation coefficient to investigate the overlap between academic and social influence and found no linear correlation. (Li and Gillet, 2013)

Both measures can be used as input to investigate different aspects of social capital. The academic influence stems from scientific expertise and highly recognized papers and is, therefore, relevant for the expertise and knowledge aspect. Social influence is a measure for a person's connections and engagement within the community. Therefore, one may be able to use a combination of both aspects as input for the assessment of a scientist's CSC. This could be achieved by averaging over both influence values and then assigning a CSC score, or by using the metrics (number of readers, R-index, centrality measures, etc.) as input for supervised learning (similar to (Su et al., 2012)).

10.4.5 *Centrality Measures in Scientific Networks*

Another article that is dedicated to the analysis of scientific networks was published by Kas et al. in 2012 (Kas et al., 2012). They investigated the change of scientific networks over time and in this context also discussed measures that are of relevance for social capital analysis.

In order to identify the key authors in the field of high-energy physics, Kas et al. used centrality measures similar to those discussed in the previous section (compare (Li and Gillet, 2013) and (Sun, 2011)). They argue that degree centrality can be used to identify intelligence, closeness centrality to a source to transmit/acquire information, and betweenness centrality as a measure to identify a person who connects groups.

For the analysis Kas et al. used a publicly available data set compiled by arXiv with 29,555 papers in the field of high-energy physics.

Closeness, and betweenness centrality were calculated on the co-authorship network degree centrality on the publication, co-authorship and citation graph. The networks used for this analysis are comparable to the multi-type citation networks discussed by Yang et al. (Yang et al., 2010).

There is no comparison to ground-truth data but the overlap between the publication out-degree and co-authorship degree is quite high, which is in line with other investigations in the field (Kas et al., 2012). The citation in-degree is interpreted as the prestige of an author and Kas et al. do not observe a significant overlap between this ranking and the other centralities. (Kas et al., 2012)

For the extraction of CSC these differences indicate that a single centrality measure is likely not sufficient to assign a CSC score. However, they can be used as input for further investigations, e.g., with supervised learning.

10.4.6 *Summary: CSC in Citation Networks*

The key takeaways of this section in the context of social capital retrieval, are:

- There are several indices (h-index, r-index, i10 index) that can be directly used to assess a person's scientific influence. (Hirsch, 2005), (Egghe, 2006) The h-index correlates with a user's social capital (Abbasi et al., 2014).
- There is a variety of different features (maximum number of readers, publications, citations, R-Index, h-index) that can be used as input for algorithms to determine someone's expertise (Li and Gillet, 2013), (Su et al., 2012).
- Different types of PageRank (homogenous or topical) can be implemented to determine a participant's influence (Yang et al., 2010).
- Other centrality measures can also be used to assess a person's influence within the network (Li and Gillet, 2013), (Kas et al., 2012).

The extracted expertise, influence, and importance of a person, which is often compiled in the form of a ranking of the scientists, can be used as input to determine CSC scores. If one is interested in dividing social capital into different topics, one can use topic models (Su et al., 2012) or leverage the fact that scientists usually publish in well defined fields with little overlap (Li and Gillet, 2013). We used the features identified in this section together with additional features as input for the CSC prediction with supervised learning in the experiment described in section 12.7.

10.5 THREADED DISCUSSION BOARDS AND Q&A PLATFORMS

Threaded discussion boards are among the most popular websites in the world and are consequently interesting for CSC assessments from online data sources. Along with previous work about the two discussion boards Reddit and Slashdot, we also discuss publications about online question and answer portals. Those are generally organized along various categories and allow users to discuss problems in a similar fashion.

10.5.1 *Reddit's Feedback Mechanism*

The scientific research about Reddit so far was mainly focused on the existing voting mechanism, which was described in section 3.5. Gilbert illustrated one limitation of the current mechanism in (Gilbert, 2013). Even though it is the purpose of Reddit to promote new and interesting ideas, Gilbert demonstrated that more than 50 percent of links that later turned out to be popular, were overlooked on their first submission. Richterich makes another interesting observation: "Karma functions as main, quasi-monetary incentive and reward of participation" (Richterich, 2013). Even though this reward has no value in the sense of real currency, many participants strive to maximize their points, e.g., by stealing content and claiming it as their own. (Richterich, 2013)

The fact that Reddit is a large, openly available collection of topic-centered discussions, makes it well suited for further investigations. A straightforward way to improve the existing metrics could be to

- separate all comments into different domains (topic distillation), and

- normalize the obtained Karma per comment (Karma/Comment ratio).

Considering the number of views per page (similar to the method presented by Gilbert (Gilbert, 2013)), the number of responses to a comment, or the number of reposts and references a popular comment creates is another way to create more features for contributive social capital analysis. These features can then be used in supervised learning algorithms to predict CSC values.

10.5.2 *Analyzing User Properties on Slashdot*

In the book "Computing with Social Trust", J. Golbeck reviewed previous work on trust in online communities (Golbeck, 2009). As there are millions of users who post daily on Slashdot.org, the author concludes that "it is often a good strategy to delegate the quality assessment task to the users themselves" (Golbeck, 2009). This is achieved by the intrinsic ranking scheme where users can award others with 'karma', similar to Reddit's approach. For the assessment of social capital this has the same consequences: an algorithm that extracts CSC scores can leverage these intrinsic karma scores.

To assess the applicability of other network analysis metrics for threaded discussion boards, one can leverage the findings of Gómez et al., who analyzed the social network and discussion threads in Slashdot in (Gómez et al., 2008).

For their investigations Gómez et al. created a graph based on the relations between authors of posts and the people who reply. They discarded several posts (e.g., anonymous ones or posts that did not inspire a discussion) and ended up with a network of more than 80,000 people and almost 1.3 million comments.

Their key findings that can facilitate future social capital research are (Gómez et al., 2008):

- The Slashdot network displays features of traditional social networks, for example large components (disconnected sub-graphs within the network), small average path length, and high clustering.
- The differences between Slashdot and traditional networks are a lower reciprocity (users do not reply as often as observed in other social networks) and neutral assortative mixing by degree (highly connected users do not prefer the exchange with other highly connected users which is often observed in other networks)
- Connections between users are less explicit as in other social networks, where they can be made public as "friendships".
- Two classes of users exist that can be separated by the mean score of their posts. Skilled writers achieve on average higher scores and run-of-the-mill writers lower scores.

The similarities between the Slashdot network and other social networks suggest that similar analysis methods can be implemented. The fact that different user classes exist additionally supports the idea of underlying CSC values that vary from user to user.

10.5.3 *Trust and Reputation on Slashdot*

In their 2009 paper Skopik et al. focus on trust and reputation in professional virtual communities, which include discussion forums like Slashdot (Skopik et al., 2009). They presented a system which determines trust relationships between community members automatically and objectively by mining communication data. This is relevant for the trust component of CSC and addresses the problem that human based ratings can be unfair and biased.

The algorithm assigns a value to comments that attract other comments, depending on the number of answers a post inspires. This algorithm was extended to model trust by implementing two functions. The first function represents the confidence of user i in user j. Skopik et al. set this function equal to the discussion strength between both users in a specific topic. The next function models the reliability. This function is set to 1 if user i interacts with user j at least 10 times within a year on a specific topic. The trust between both users was then calculated as a product between confidence and reliability function. (Skopik et al., 2009)

The experimental validation was conducted on data that was crawled from two Slashdot subdomains over the course of one and a half years. After the data were cleaned (e.g., subtracting anonymous users) the graph had 24,824 nodes and 343,669 edges. The sets of trusted users identified by the algorithm were similar to the sets identified by humans. (Skopik et al., 2009)

It is possible to refine the results, for example, by implementing more complex functions for confidence and reliability. The calculated trust can be used as input for CSC extraction.

10.5.4 *Authority Identification in Q&A Portals*

In their 2015 paper, Bouguessa et al. propose methods to identify authoritative actors in online communities, like question and answer portals. As Q&A portals are content focused, it is reasonable to assume that helpful contributors achieve a higher authority score than others. Consequently, their contributive social capital should be higher — which underlines the relevance for this paper.

Bouguessa et al.'s approach can be divided into two steps. At first, a feature vector is derived for each user that contains information on the user's social activities. Then, a statistical framework, based on the multivariate beta mixtures, is used to model the estimated set of feature vectors. This is used to identify the most authoritative users in the network. (Bouguessa and Ben Romdhane, 2015)

On the question and answer platform Stack Overflow, Bouguessa et al. used the following features as input:

- Number of answers,
- number of answers that were chosen as "best answer" by the community,
- number of votes received, and the
- Z-score $= \frac{a-q}{\sqrt{a+q}}$, where a (q) is number of answers (questions) by the user.

To best distinguish between regular and important users, these features are transformed to a logarithmic scale. This compresses the long tail of higher values that correspond to authoritative users, while stretching out the smaller region which corresponds to the large number of less influential users (e.g., low count of best answers). Additionally, all features are normalized into the interval $[0, 1]$.

Bouguessa et al. propose that these normalized user feature vectors follow a mixture density. Its parameters can be determined with maximum likelihood estimation. Authoritative users can then automatically be determined by choosing people whose feature values lie above a certain threshold.

A verification of this algorithm on Stack Overflow and Twitter data yielded high quality results (Bouguessa and Ben Romdhane, 2015).

The features, as well as the multivariate approach, can be used for CSC prediction, e.g., with supervised learning.

10.5.5 *Summary: CSC on Threaded Discussion Boards*

The most relevant findings for CSC extraction from threaded discussion boards and question and answer portals are summarized as follows.

- Richterich and Gilbert discussed the current rating algorithms of Reddit (Richterich, 2013), (Gilbert, 2013). Based on their findings we concluded that the intrinsic evaluation tool ("Karma points") should be investigated for social capital extraction and propose to implement ratios for normalization as we have seen for Twitter (compare Anger et al. (Anger and Kittl, 2011)).

- Similarly, for the online community Slashdot, J. Golbeck suggests the use of the internal quality assessment mechanisms to evaluate user's trust (Golbeck, 2009) — a constituent of CSC.

- There are similarities between Slashdot and other social networks (Gómez et al., 2008) that suggest that similar analysis methods can be used to investigate user's features and consequently social capital.

- The features on question and answer portals can be used to identify authoritative users (Bouguessa and Ben Romdhane, 2015) and likely also used for CSC prediction.

We investigated CSC extraction from threaded discussion boards and Q&A portals in an experiment on Quora. This experiment is discussed in section 12.8.

11

INVESTIGATION OF CSC EXTRACTION FROM OSN

The first section of this chapter briefly summarizes the experiment and the results (secion 11.1). In section 11.2, the motivation is presented and the investigated research questions are introduced. Then the experiment (section 11.3) and the created data sets (section 11.4) are discussed. This discussion also entails a description of the participants, the collected features, and the ground truth assessment. The analysis is described in section 11.5. Potential shortcomings of the experiment are discussed in section 11.6, before section 11.7 discusses the results with regard to the initial research questions. Section 11.8 sets the findings in context to the overarching research questions of this thesis.

11.1 SYNOPSIS

The first pillar of the CSC system plans to leverage OSNEM analysis to evaluate the CSC of people. In the preceding chapter, previous work about the extraction of characteristics related to social capital from different online data sources was reviewed. We have seen that supervised learning may potentially be a suitable approach to extract personal expertise, influence, reputation, and trust. We reasoned that because of the relation of these terms with CSC, it should be possible to extract this as well. To investigate this hypothesis, we conducted an experiment with 242 participants. In the course of this thesis the experiment is referenced as *social capital experiment*.

A data set was collected with the help of a dedicated social networking platform to which all participants had access at any time. It combined features from the popular networking platforms Facebook and Twitter. People could interact with each other in the form of blog-posts, comments, and messages, follow each other, and "like" helpful contributions. Based on this data and additional ground truth approximations collected with questionnaires, several machine learning algorithms were compared. Most algorithms were better than a random predictor, which indicates that CSC extraction from social networks may be possible. However, the algorithms do not allow to predict the CSCW of each participant correctly, which limits the applicability of the predicted CSCWs.

11.2 MOTIVATION AND RESEARCH QUESTIONS

Being able to extract CSC values from previous interactions, i.e. contributions by the users and feedback from others, on social media is an important part of the first pillar of the CSC system. This allows to leverage previous contributions and micro assessments (e.g., likes) of and about users without requiring any additional work of the users. This ease of operation is important to motivate people to participate in the system. The main research questions in this context are:

1. Can CSC scores be predicted for each participant based on previous network interactions?

2. Which supervised learning algorithm is best suited for the prediction?
3. Does the activity of the users play a role for the accuracy of the prediction?
4. Which features are most important for the prediction?

11.3 DESCRIPTION OF THE EXPERIMENT

The study was conducted as the practical part of a course on social computing in the summer term 2017 at Technical University of Munich in a way that combined the goals of the lecture with the goals of the experiment. The goals of the tutorial of the lecture were:

- Convey social computing and network analysis concepts (python programming, centrality measures, network analysis, social computing related clustering, and other data-mining techniques).
- Create a pleasant learning environment for the students.

And the goals of the experiment were:

- Create a real, large, and lively social network for a couple of weeks with contributions from many participants that can be used for the feature extraction.
- Collect ground truth data that approximates the true CSC of a person with the help of peer assessments.
- Motivate people to participate in a natural way, without incentivizing fake or biased data.
- Secure the students' privacy.

The key to combining these goals was to find an application for the data sets required for the experiments that also allowed the students to work with them. This was achieved by creating exercise sheets that included, among others, network analysis exercises. The created social networking data could then be used as input for the calculations as well as later for this experiment.

For the creation of the data sets, several methods were leveraged that have been suggested for the enhancement of student experiences in massive courses (Bry et al., 2014): The students could use the social networking platform for discussions about the lecture and other topics. Dedicated online questionnaires acted as response systems and allowed the students to provide feedback about their experiences. Most importantly, they could directly work on anonymized excerpts from the social networking platform to test algorithms learned in the lecture.

According to individual feedback and the discussions on the social networking platform, the students perceived this as generally more exciting than analyzing artificial social networks.

To prevent fake data, the participation in the network was voluntary and to guarantee privacy, the students were free to choose pseudonyms. A screenshot of the web interface of the social networking platform is provided in figure 7.

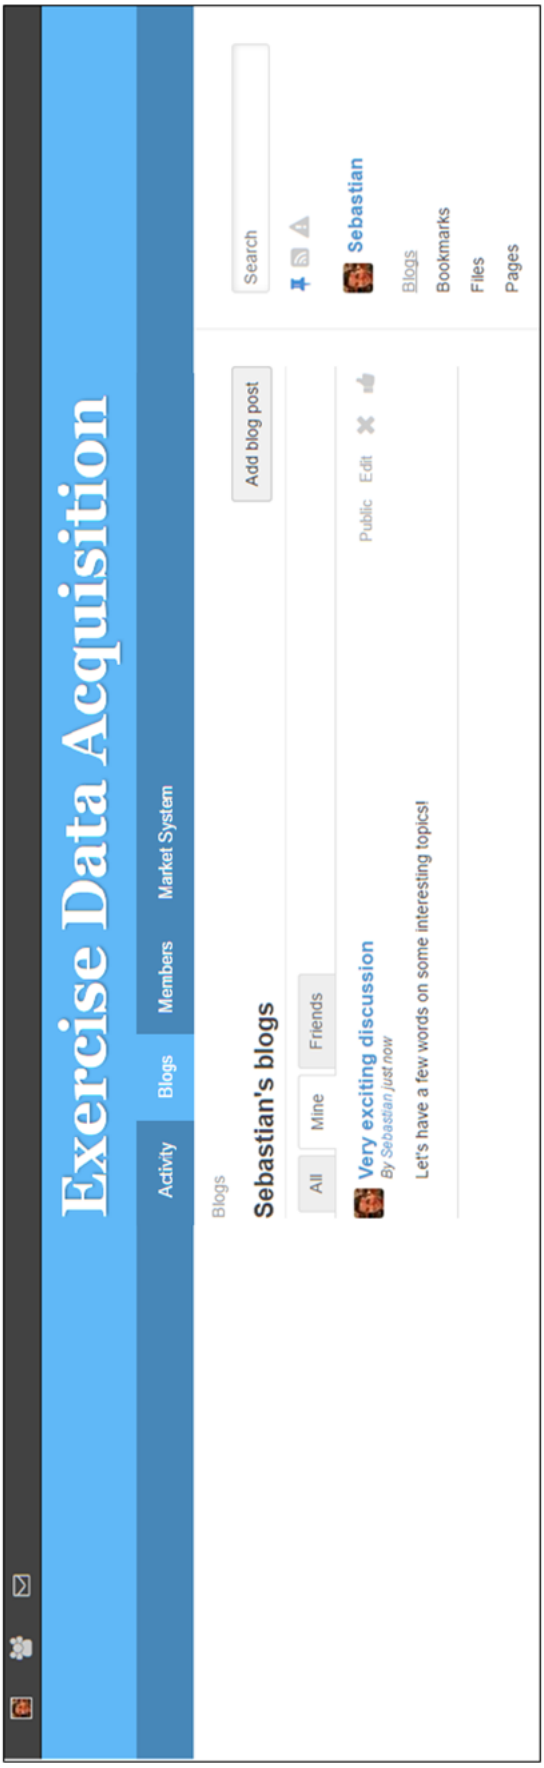

Figure 7: Screenshot of the social networking interface of the social capital experiment

The platform title **Exercise Data Acquisition** indicates that the students would later work on their own data in the exercises. Below are four buttons. **Activity** leads to a page that displays the recent activity (e.g., blog posts, new friendships, etc.). **Blogs** leads to a list of all discussion posts, which users can comment on or "like". **Members** displays a list of all participants, including their profile pictures. Users can browse this section to start friendships or send private messages. The button **Market System** is a direct link to the social capital market, in which the participants could trade currency. The market interface is illustrated in figure 27.

Additional motivation to participate was provided in the form of conversation starters during the weekly discussion of the exercises. The topics were of interest to the students and ranged from current events to the best restaurants in Munich. Some exemplary post headlines are listed in appendix B. Of over 400 students who took the course, 242 registered to the system and for 165 we collected at least one ground truth assessment by others. The networking platform was based on Elgg[1], an open source framework for creating custom social networking platforms. Users were provided with a functionality similar to Facebook and Twitter. They could create profiles with pictures, follow one another, write posts, or comment on own or other people's posts. They could also "like" posts and comments and send private messages.

Students could contribute to the social networking platform during a timespan of nine weeks in the middle of the semester. During this time the social networking platform and the market system (which is described in detail in chapter 14) could be accessed at any time. There were six exercise sheets that were presented and discussed in seven tutorials. The content of the exercise sheets was as follows:

- **Exercise sheet 1:** Python programming and introduction to graphs (Exercises: Use pyhton to write a Fibonacci series; Write a "guess my number" game; Create an object that represents the Krackhardt kite graph (Krackhardt, 1990); Create a Watts-Strogatz small world graph (Watts and Strogatz, 1998))

- **Exercise sheet 2:** Introduction to centrality measures (Exercises: Calculate the degree centrality of the Krackhardt kite graph; Calculate and compare the degree centralities of the course participants on the friendship, comment, and like graphs extracted from the social network (anonymized); Calculate the closeness centrality for each node on the weighted like graph (anonymized); Calculate and interpret the betweenness centrality of each node on the friendship graph (anonymized))

- **Exercise sheet 3:** Introduction to social network analysis and social capital prediction (Exercises: Investigate correlations between features extracted from the network and anonymized ground truth assessments with scatter plots; Apply logarithmic transformation to the features; Predict CSC scores with linear regression; Interpret the results with QQ-plots and residual analysis)

- **Exercise sheet 4:** Introduction to graph clustering (Exercises: Cluster the course network (anonymized) with k-means clustering; Cluster the Krackhardt kite graph with the Girvan Newman algorithm)

1 https://elgg.org/

- **Exercise sheet 5:** Working with the social capital market and introduction to database analysis (Exercises: Translate SQL statements into pandas; Query market participants with specific properties from the social capital market (anonymized); Calculate different transaction ratios of participants of the social capital market (anonymized))

- **Exercise sheet 6:** Fraud detection and repetition of previous methods (Exercises: Identify market participants with high CSCW (anonymized); Identify the transaction partners of the participants with high CSCW (anonymized); Compare the results with the relationship graph of the social network; Compare the results to the interaction patterns of the social network; Detect fraudulent behavior)

11.4 DESCRIPTION OF THE CREATED DATA SET

During the experiment 244 posts, 2,868 comments, 1,930 following relationships, and 3,651 likes were created by the 165 participants for whom a ground truth assessment was available. Users were free to write about what they wanted. To encourage discussions, different conversation starters in controversial topics were provided in the lecture: populism in politics, living in Munich, and healthy food and sustainability. The topics were chosen with regard to two main concerns. They should motivate students to contribute and should be controversial in a way that encouraged discussions between different opinions. All three topics were actively discussed by the users. The persons conducting the experiment did not take part in the experiment and did not review or comment on any discussions in order not to influence the dynamics of the network. One teaching assistant was responsible for monitoring the network for obscenities, insults or any other unwanted or illegal behavior.

The observed behavior was similar to what one might see on platforms like Facebook or Twitter. Some people discussed current affairs, others posted funny content, memes, and sometimes spam-like messages or advertisements (e.g., for university events that they organized).

11.4.1 *Ground Truth Assessment*

To approximate a ground truth value, the students were asked to answer a questionnaire about other students in the course. They were presented with a list of all students who registered and could select as many as they felt confident to assess. There was a total number of 539 estimations for 165 students, an average of 3.3 assessments per person. Only these 165 people were considered for further analysis. This approach is similar to the ground truth peer assessment used in previous work (Campbell et al., 2003).

The questionnaire consisted of 8 questions, each associated with one of the three hypothesized contributing factors of CSC: competence, trustworthiness, and social responsibility. All assessments were made on a scale with 100 unmarked steps.

- The competence assessment should be related to the knowledge and expertise a person demonstrated in the network. Therefore, we asked for direct assessments in the three topics that were provided as discussion starters during the experiment. For these questions the left side (0) of the scale was labeled "no experience at all", the right end (100) "extremely knowledgeable".

- The trust assessment was supposed to assess to what degree the individual was trusted by others. Three questions were used that were inspired by the research by Jones et al. on diagnosing trust (Jones and Shah, 2015). They elicited, with an overall assessment of trust, the belief that the other person was concerned with the other's welfare, and finally the feeling to what extent the person is fair and honest.
- The third part of CSC, the social responsibility, was assessed with two questions. The first asked about the environmental friendliness of the person, the second about their level of social support and engagement.

The eight questions provided the participants with a multi-faceted way to assess their counterpart, with the individual characteristics being easier to assess than contributive social capital directly. The full questionnaire is provided in appendix A. A single CSC value per person was calculated by averaging over all values. This score represents the perception of an individual by that person's peers along the three dimensions. In the absence of a more suitable CSC value, this assessment was used as approximation for the ground truth for the following analyses. The mean CSC value obtained this way was 64.0 with a standard deviation of 11.5, minimum value was 29.8, and maximum value 94.8. The distribution is visualized in figure 8. The Shapiro-Wilk-Normality test reveals a likelihood of 0.935 with a p-value < 0.0001 that the ground truth is normally distributed.

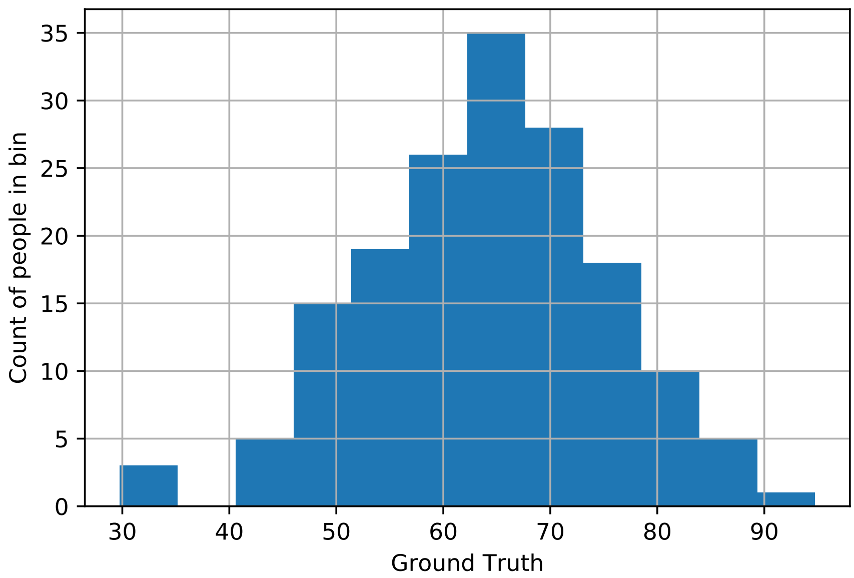

Figure 8: Distribution of ground truth assessments in the social capital experiment

Feature	Mean	σ	Min	Max
Posts	1.5	3.0	0	24
Comments	17.4	39.6	0	47
Liked posts (active)	6.2	8.1	0	47
Liked comments (active)	15.9	49.1	0	581
Liked posts (passive)	4.3	7.3	0	34
Liked comments (passive)	12.7	35.1	0	353
Comment responses to posts	16.4	30.5	0	176
Messages sent	4.2	30.3	0	384
Messages received	3.1	6.6	0	68
Followers	11.7	12.0	0	104
Friends	13.6	31.4	0	347
Characters posts	350.5	734.8	0	5,331
Characters comments	1,790.5	3,578.8	0	34,696
Characters messages	348.5	2,458.9	0	30,128

Table 1: Collected features from the social networking platform, including the mean count, the standard deviation, and minimum and maximum values.

11.4.2 *Demographics of the Participants*

76.4% of the 165 students were male, 23.6% female. The average age was 23.2 years. 35.2% were between 18–21, 43.0% between 22–25, 15.8% between 26–29, and 3.6% between 30–35. 2.4% decided not to disclose their age. The nationality was mainly German (64.2%), 8.5% were from India, 2.4% from Turkey, and the remaining 24.8% from 27 other countries.

11.4.3 *Interaction Features in the Network*

The average number of contributions to the OSN by these 165 people is illustrated in table 1. It lists the respective features, their mean value (e.g., number of posts per person), the standard deviation, as well as the minimum and maximum values.

For all collected features a small number of people is responsible for the majority of the contributions, which resembles a power-law distribution. This is visualized in figures 9, 10, and 11 for the number of written comments, followers, and the number of likes received on comments. This is in line with what we see in larger networks (Rastogi, 2016), (Gjoka et al., 2009).

11.5 csc analysis

The main purpose of our analysis was to investigate the research question "can a person's contributive social capital be approximated based on their interactions in a

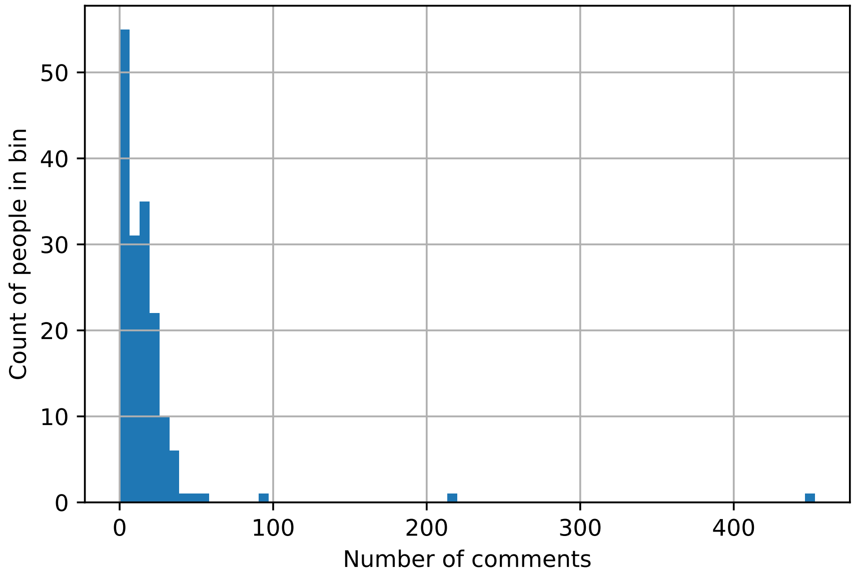

Figure 9: Histogram of the comment contributions to the network in the social capital experiment

social networking platform?" For this purpose we ran two different analyses on the whole data set and an active user subset.

11.5.1 *Preprocessing*

The first step of the analysis was to create a data set of people from whom we have a ground truth approximation, as well as a registration to the networking platform. As ground truth we chose the assessment by the person's peers, as described in section 11.4.1.

All features that characterize the activity, interactions, and feedback on the social networking platform were extracted with a Lenovo Yoga 510 laptop with 8 GB RAM. A list of all features used for the analysis is provided in table 1.

The resulting data set included 14 features and one ground truth CSC approximation for 165 participants. This data set was saved in a database and used for the following analyses.

11.5.2 *Prediction and Correlation with the Whole Data Set*

Two different methods were used to test the hypothesis. The first evaluation was to predict contributive social capital scores based on the features extracted from the network (see table 1) in comparison to a baseline estimator. The group of users was ranked according to their predicted CSC score and then compared to a ranking based on the ground truth assessments.

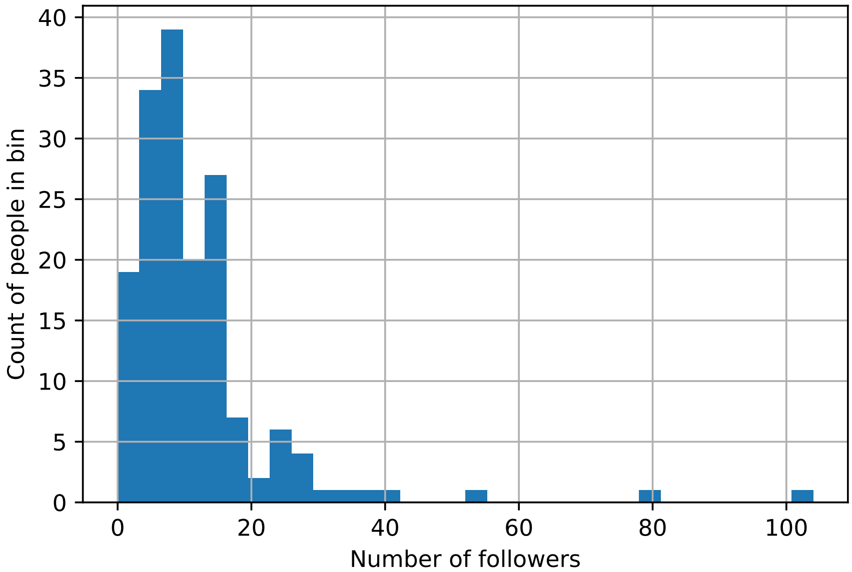

Figure 10: Histogram of the number of followers in the social capital experiment

11.5.2.1 *Prediction of CSC Scores Based on Network Features*

Several different algorithms were used for the evaluation: linear regression (with and without regularization directly using the features listed in table 1), as well as regression with a decision tree, a random forest, and a neural network. We used 10-fold cross validation for all algorithms. The neural network had 200 neurons in one hidden layer and a logistic sigmoid function as activation function for the hidden layer. The random forest regressors had ten trees and no restrictions on the maximum depth of the tree. To evaluate the result, the mean average error of each model's predictions (mean difference between predicted social capital score and the ground truth approximation) was compared to a baseline predictor that always predicts the mean ground truth approximation of the training data. The results are summarized in table 2. To evaluate the quality of the fit, residual plots and QQ-plots were created. The residuals were distributed evenly and the QQ-plots did not exhibit any systematic deviations from the expected straight line.

The best result is achieved with random forest regression, which performs almost 17 percent better than the baseline predictor. This is followed by a decision tree with depth four. Linear regression with Lasso regularization, the neural network, and linear regression are only marginally better than the baseline predictor.

11.5.2.2 *Ranking of People Based on Their CSC*

For the ranking task we used the same algorithms to predict a CSC score for each user. All participants were then ranked according to the predicted value. The correlation between this ranking and a ranking with the ground truth approximation was used to evaluate the goodness of the prediction. The results are summarized in table 3.

Algorithm	**Mean absolute error**	**Improvement**
Baseline	9.11	–
Linear regression	9.03	0.8%
Linear regression with regularization	8.73	4.2%
Decision tree	8.45	7.2%
Random forest	7.57	16.9%
Neural network	8.87	2.6%

Table 2: Performance of the different algorithms compared to a baseline predictor for all 165 users. The improvement indicates by how much the algorithm outperforms the baseline.

Algorithm	**Pearson**	**Spearman**
Linear regression	0.24 (0.0019)	0.44 ($<$ 0.0001)
Linear regression with regularization	0.29 (0.0001)	0.46 ($<$0.0001)
Decision tree	0.44 ($<$ 0.0001)	0.41 ($<$ 0.0001)
Random forest	0.42 ($<$ 0.0001)	0.41 ($<$ 0.0001)
Neural network	0.29 (0.0002)	0.29 (0.0002)

Table 3: Pearson and Spearman correlation of the respective algorithm between the predicted ranking and the ranking according to the ground truth assessments. The first value is the correlation, the value in brackets the p-value.

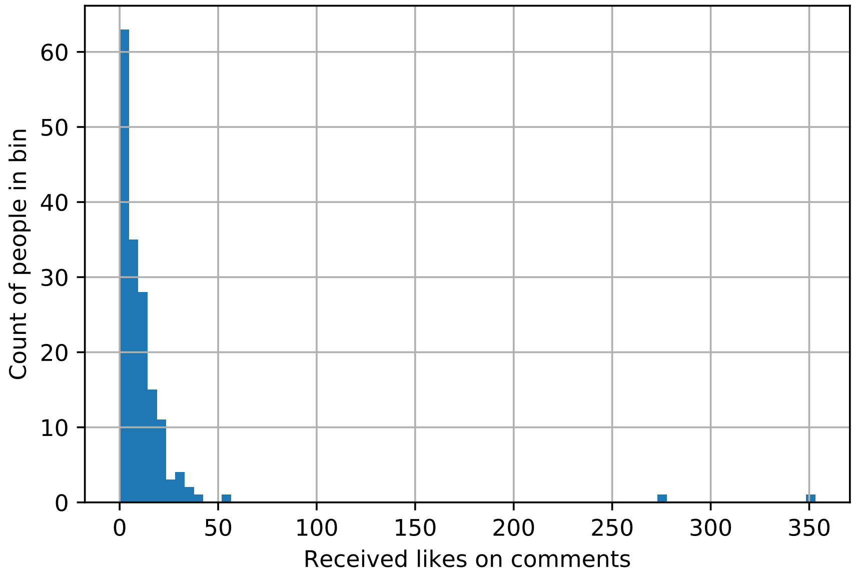

Figure 11: Histogram of the number of likes participants in the network received on their comments in the social capital experiment

For all algorithms we can observe a weak to moderate positive correlation. The p-value indicates a statistical significance at the 0.01 level for all algorithms. That means the probability is less than 1 percent that these or larger correlation values are observed if the null hypothesis – that there is no correlation between the ground truth and the features – is true. The largest Pearson correlation was achieved with the algorithms decision tree ($r = 0.44$) and random forest ($r = 0.42$) that both achieved the second best and best improvements in the previous analysis. The highest Spearman correlation was achieved with regularized linear regression ($\rho = 0.46$). Decision tree and random forest regression also ranked the users in a way that correlates moderately with the ground truth ranking ($\rho = 0.41$ and $\rho = 0.41$). The case of linear regression and neural network regression is particularly interesting. Both algorithms only showed marginal improvements in the first analysis. In the ranking they demonstrated a weak positive correlation of $r = 0.24$ for linear regression and $r = 0.29$ for the neural network. This might indicate that it is easier for algorithms to rank people according to their CSC than to predict a concrete value. These results are discussed in section 11.7.

11.5.3 *Prediction and Correlation with an Active User Subset*

Some of the students in the data set of 165 participants contributed little to the social network. Therefore, we ran a second analysis with only active members to investigate potential differences. We performed the same two evaluations as in the previous section, this time only with users who wrote at least one post or comment and who did befriend at least one other user. This data set contained 139 active participants. The residual plots and QQ-plots did not exhibit any particularities.

Algorithm	Mean absolute error	Improvement
Baseline	9.06	–
Linear regression	9.86	-8.8%
Linear regression with regularization	8.98	0.9%
Decision tree	7.94	12.4%
Random forest	7.20	20.6%
Neural network	7.27	19.7%

Table 4: Performance of the different algorithms compared to a baseline predictor for the subset of 139 active users. The improvement indicates by how much the algorithm outperforms the baseline.

11.5.3.1 *Prediction of CSC Scores Based on Network Features*

For the group of active participants the average ground truth CSC assessment was 65.0 and, therefore, marginally higher than the whole group's value of 64.0. All parameter settings of the employed algorithms were the same. The results of the prediction are summarized in table 4. These results are discussed in section 11.7.

The result is considerably better than for the whole data set. Both, random forest regression (20.6%) and neural network regression (19.7%) led to an improvement of about 20 percent. The decision tree also demonstrated slight improvements (12.4%). The simple linear regression algorithm performed worse than on the full data set and has a larger mean error than the baseline predictor. With Lasso regularization it performed only marginally better. This indicates that the relation between the network features and the ground truth CSC assessment can be described less well by a linear function when inactive users are excluded.

11.5.3.2 *Ranking of People Based on Their CSC*

The ranking of students according to their predicted CSC scores reflects a similar result, as illustrated in table 5.

All algorithms ranked the active users in a way that their CSC rank correlates positively with the ground truth value. The highest Pearson correlation ($r = 0.63$) was achieved with the neural network. Random forest and decision tree ranking also achieve values larger than 0.5 ($r = 0.59$ and $r = 0.53$). The relatively weak correlation of $r = 0.22$ that was achieved with the linear regression indicates once more that the relation between ground truth and features is most likely not purely linear. When using Spearman correlation, the best results are also achieved with random forest ($\rho = 0.49$) and the neural network ($\rho = 0.47$). However, all correlation values lie much closer together.

11.6 POTENTIAL SHORTCOMINGS OF THE EXPERIMENT

There are several potential shortcomings of the data set that need to be mentioned.

Algorithm	Pearson	Spearman
Linear regression	0.22 (0.0079)	0.43 (<0.0001)
Linear regression with regularization	0.34 (<0.0001)	0.46 (<0.0001)
Decision tree	0.53 (<0.0001)	0.46 (<0.0001)
Random forest	0.59 <0.0001)	0.49 (<0.0001)
Neural network	0.63 (<0.0001)	0.47 (<0.0001)

Table 5: Pearson and Spearman correlation of the respective algorithm between the predicted ranking and the ground truth ranking. The first value is the correlation, the value in brackets the p-value.

- The majority of contributors are university educated, male students between 20 and 30 years of age. This is not representative of the total population, even though, the users of social networking platforms are predominantly below 35 and more often male (Statista, 2017).

- The sample size of 165 is relatively small.

- The cross-over of data collection and university lecture might have led to a bias. The participation in the experiment was voluntary and did not influence the grade in any way. However, we cannot exclude that some students only participated or contributed in a certain way because they hoped to create a positive impression. We tried to counter this bias with open and generous discussion culture during the experiment.

- A similar bias is possible regarding the ground truth assessment. The assessments were confidential and were never displayed to the users. However, there might be a positive bias because students might not have wanted to assess others negatively. This kind of bias can be expected in all experiments of this type.

- The time frame of nine weeks is short compared to other OSNEM that are running for years. This led to scarce data for some less active participants. For those individuals it was difficult to predict correct CSC values.

- We chose an averaged peer assessment of each individual along eight questions as ground truth approximation. The eight questions reflected the three dimensions of CSC, as discussed in section 2.1.2. While this is an external assessment that is likely to correlate with or even represent the true CSC of a person to some extend, it was not possible to assess the overlap.

All these shortcomings need to be kept in mind when interpreting the results of the study.

11.7 DISCUSSION OF RESULTS WITH REGARD TO THE RESEARCH QUESTIONS

The experiment with the active user subset yielded an improvement of about 20% over the baseline predictor when trying to predict CSC values, and a Pearson correlation value on the ranked lists of up to 0.6. These values indicate that it might be possible to predict contributive social capital from activity and feedback features present on social networking platforms (research question 1). However, the values are merely small to moderate, which indicates that there are limitations for assessments purely based on the analysis of features extracted from OSN. Additionally, the calculated results might be biased due to the shortcomings of the experiment, as we discussed in section 11.6. Therefore, it is important to use caution until the findings are supported by large-scale experiments with data from existing social networking platforms. We are not aware of any similar experiments for the analysis of contributive social capital in social networks, therefore, it is hard to compare the values of our results. Nevertheless, one can make several other observations:

- The best algorithms for the prediction of contributive social capital are random forest (best results for both data sets) and the neural network that performed also well on the active user network (research question 2). Support vector regression (SVR) was also tested but did not perform better than any of these algorithms. The comparatively unsatisfying performance of multiple linear regression indicates that linear correlation does not sufficiently describe the relationship between the ground truth assessment and the respective features. The non-linear relationship can be learned by both random forest and neural network. The neural network can express the non-linearity directly due to its non-linear activation function. The random forest describes the non-linearity with a piecewise linear approximation to describe the relation between input and output. The approximation makes the random forest more parsimonious regarding the number of parameters and thus more easily interpretable.

- On the active user data set, the best results are obtained with the neural network and random forest (Pearson correlation), respectively with random forest and the neural network (Spearman correlation). For the ranking on the whole data, the best result is achieved with decision tree and random forest regression (Pearson correlation), or linear regression with and without regularization (Spearman correlation). The results are consistent with the previously established performance of the algorithms. The high Spearman correlation on the whole data set that was achieved with the linear regression ranking may be explained by the gap between the linear model and the non-linear data, which results in a non-linear error term.

- It appears that a better result can be achieved by ranking the users than by predicting concrete CSC values.

- The quality of the analysis was increased by excluding inactive users (research question 3). This led to an improvement of 20.6% compared to 16.9% for the prediction task, and a Pearson correlation of 0.63 as opposed to 0.44.

Feature	**Importance**	**Cumulative**
Liked comments (passive)	24.9%	24.9%
Comments	15.0%	39.9%
Characters posts	14.2%	54.1%
Comment responses to posts	11.0%	65.1%
Followers	7.5%	72.6%
Characters comments	5.2%	77.7%
Liked comments (active)	4.9%	82.7%
Friends	4.4%	87.0%
Characters messages	2.9%	89.9%
Liked posts (active)	2.5%	92.4%
Messages received	2.3%	94.7%
Messages sent	2.1%	96.8%
Posts	2.1%	98.9%
Liked posts (passive)	1.1%	100.0%

Table 6: Relative and cumulative importance of the different features for random forest regression on the active user data set

The importance of the different features for the prediction can also be investigated. As random forest regression generally led to the best results, we chose this algorithm to discuss their relative importance. As illustrated in table 6, the five most important features are the number of likes a user received on their comments, the number of comments written by a user, the number of characters used in written posts, the number of comments that a post inspired, and the number of followers a user has (research question 4).

These five features account for over 70% of the importance for the model (increase in prediction error when leaving out the feature (Breiman, 2001)). Three of the features are indicators for the support a user receives from their surrounding network (feedback features), namely the received likes, the inspired responses, and the number of followers. The number of comments as well as the length of the posts are signs for the involvement of a user (activity features). Other ways of participating, like following others or liking the posts or comments of other users are less important for the prediction of CSC. The number of posts and the likes received on them are on the bottom of the list. This might be due to their relatively low number (on average 1.5 for posts and 4.3 for likes on posts per person). A variety of additional features, including those unique to different data sources, were investigated in the experiments in chapter 12.

11.8 SUMMARY AND OUTLOOK

This chapter presented an experiment that was conducted to investigate contributive social capital in social networks. In its course a data set with network activity and ground truth CSC assessments for 165 participants was created. The first investigation was to predict CSC scores based on the network activity and compare them to a simple baseline predictor. The second was to rank people according to their predicted CSC values and correlate the result to the true ranking. There was a small improvement regarding the prediction and a moderate correlation between both lists. However, this is just a piece of evidence for the predictability of CSC in OSNEM and not a definite proof due to the limitations of the experiment.

Nevertheless, the experiment yielded important insights for the CSC system. The positive correlations and improvements over the random predictor indicate that the first pillar of the CSC system can be an important part of an assessment. This also contributes to the third overarching research question introduced in section 1.2, whether such an assessment is possible. The results are in line with previous research about the extraction of other CSC related properties from online social networks (e.g., (Ziegler, 2009), (Yang et al., 2010), and (Su et al., 2012)). The limited accuracy of the prediction underlines the importance of additional assessments that the CSC system suggests in the form of market interactions as well as certifications and endorsements.

To address the shortcomings of the experiment, which were discussed in section 11.6, large scale investigations are required — ideally in a natural environment over the course of several years. The following chapter describes four experiments we conducted about the analysis of CSC from Facebook, Twitter, Quora, and on a scientometrics data set.

12

INVESTIGATION OF CSC EXTRACTION FROM ONLINE DATA SOURCES

The following experiments were conducted in the context of the master's theses by Monika Varshney (Varshney, 2017), Valeriia Chernenko (Chernenko, 2017), Johannes Feil (Feil, 2017), and Rauf Zeynalov (Zeynalov, 2018), all of which were supervised by Sebastian Schams and Georg Groh in the Social Computing Research Group at Technical University Munich.

This chapter is structured as follows. After an initial summary of the experiments (section 12.1), a problem motivation and the introduction of the research questions (section 12.2), as well as a privacy disclaimer (section 12.3), the important question of how to find a suitable representation of CSC that can be used to approximate ground truth values, is discussed (section 12.4). The following four sections are all structured the same way and describe the CSC investigations on Facebook (section 12.5), Twitter (section 12.6), scientometrics (section 12.7), and Quora (section 12.8). All of these sections contain a discussion of the respective data set, more details on the selected ground truth approximation, the analysis approach that was tailored to the respective platform, and a discussion of the results. Finally, section 12.8.4 compares the results and sets them in context to the CSC system and the overarching research questions.

12.1 SYNOPSIS

In the previous chapter 11, we presented an experiment to investigate the extraction of CSC scores from a dedicated social networking platform. In addition to this investigation, we ran several experiments to examine the extraction of contributive social capital from four other online data sources. In chapter 3, five different data sources that contain information about CSC were presented. In this chapter, we compare methods for the extraction of CSCWs from four of these sources, a social networking platform (Facebook), a microblogging service (Twitter), scientometrics (on an ArnetMiner data set), and a Q&A/threaded discussion portal (Quora). Direct communication was omitted, as the data is usually private and consequently difficult to use for a continuous CSC assessment. If required, the following investigations and findings can be transferred to direct communication, as well.

The data used for the experiments was collected from the respective platforms and represents samples of the whole communities. The data sets for the different platforms are independent from each other. It was not possible to obtain CSC ground truth assessments with the help of questionnaires as we did in chapter 11. Consequently, suitable ground truth alternatives were used for the analysis on every platform.

The analyses were tailored to the different platforms and used different information, activity, feedback, and centrality features. The prediction was performed with different supervised learning algorithms: linear regression, linear regression with regularization, decision tree, random forest, and a neural network. Lasso was used for the regularization of the linear regression. By penalizing the absolute value of the magnitude of the coefficients, some lesser important features are set to zero, which increases the generalizability of the algorithm. The maximum depth of the decision tree and the random forest is 20. The neural network consists of 88 neurons in one hidden layer and used a logistic activation function. These parameters were determined with hyperparameter optimization with the hyperopt[1] framework (Bergstra et al., 2013).

The performance of the different algorithms was compared regarding the R^2 value, the mean absolute error, and the Pearson and Spearman correlation coefficients that were calculated between the predicted CSCWs and the respective approximate ground truths.

12.2 MOTIVATION AND RESEARCH QUESTIONS

The following experiments build upon the study presented in chapter 11 and have a similar motivation: To investigate whether CSC can be extracted from different online data sources. The specific research questions are:

1. Which data from Facebook, Twitter, Quora, and the ArminMiner data set allows an investigation of the CSC scores for each participant?
2. Which features can be defined on these different data sets to be used for the analysis?
3. What can be used as approximation for the ground truth CSC value?
4. Can CSCW be extracted based on these features with the help of supervised learning algorithms?
5. What is the best algorithm for the prediction?
6. What are the most important features for the prediction?

In the following sections 12.5, 12.6, 12.7, and 12.8, these questions are investigated for the respective platforms. A summary and discussion of all experiments with regard to the CSC system is provided in section 12.9.

12.3 DATA PRIVACY

All data that was collected for the investigations was publicly available on the respective platforms. In the recent months, public awareness regarding the use and analysis of data from social networking platforms has risen. This is mainly a result of the misuse of data from the networking platform Facebook during the 2016 American presidential election (NY-Times, 2018). We welcome this debate and ensure that all data that we obtained is treated confidentially and only used for scientific research.

1 http://hyperopt.github.io/hyperopt/ (retrieved 2018-06-15)

12.4 finding alternative ground truth approximations for csc

The CSC experiment we conducted with students at TUM university (see chapter 11), allowed us to collect ground truth assessments with the help of questionnaires. As stated in the beginning of this chapter, it was not possible to obtain a similar assessment for the network users of the experiments presented in this chapter. For each platform, we investigated a variety of alternatives. There are three main options for such a ground truth approximation. The first one is the already mentioned self or peer assessment with questionnaires. The second is the use of external ratings that correlate with the CSCWs. The third option is to use one or some of the features that were extracted for each user and use it as approximation for the full CSC value.

The first option provides a detailed assessment of each person that is likely to be the closest to a real ground truth value. Due to the time and effort required for such an assessment, we had to dismiss this option. Therefore, we closely investigated the two alternatives. Option 2 has the advantage that it leverages external knowledge that is not yet available in the network. There are several scores that may be used as ground truth approximation. We crawled, e.g., the Klout score (see also section 10.2.3) for many users of the Twitter data set and investigated it as ground truth approximation. Other investigated values were, for instance, the score the universities of the scientists in the scientometrics database obtained in rankings or the h-index of the scientists. However, all of these scores have two distinct disadvantages: They were created with a specific application in mind that is different from contributive social capital and there are assignment problems. The latter means that while a user name in the crawled data set and in the ground truth source may be similar, it does not guarantee that this is indeed the same person. The only reliable alternative for a ground truth value is consequently the use of features that were crawled when creating the data set. That means out of all features $A_1, A_2, ..., A_n$ of the data set, we select a feature B that is no longer used as input feature but as approximation of the real ground truth C. The advantage is that this score is available for every person in the data set and can be unambiguously assigned. The downside is that B is only a high level approximation. One could also argue that this makes the whole analysis redundant as one could simply use the normalized value of B for the assessment. This disregards, however, all the insights that can be created with B as ground truth:

- The value B' that was predicted by using B as ground truth and $A_1, A_2, ..., A_n$ without B as features, may be a better approximation of the real CSCW as it includes information from all features $A_1, A_2, ..., A_n$.
- The prediction of B' allows to investigate relations between the input features $A_1, A_2, ..., A_n$, e.g., to identify redundant features.
- If a value B is selected that has a relation to the real ground truth, the importance of the features for the prediction of B' is likely also relevant for the prediction of the real CSCW. These features can then be leveraged for further investigations.
- There may be situations in the future where B is not available. The investigations of this chapter can then be used to predict it.

- If the prediction of a person's CSC were purely based on an individual feature B, the system would be open to fraud and manipulation. Twitter followers, e.g., can be bought (Yang et al., 2011) which could have a massive influence on B but be of less importance for a more complex CSC prediction.
- These investigations pave the way for all future experiments where a real ground truth value is available. This real value can then directly be used instead of B — all other investigations can stay unaltered.

These arguments only hold true if B is selected in a way that ensures that it is a suitable ground truth alternative. To define such an alternative, we can leverage the results of the previous experiment that had a comprehensive ground truth assessment available. Table 6 lists the most important features for the analysis. The most important feature was the number of likes received from others and all top features were either a measure for the contributions to the network (e.g., number of comments) or for the feedback from others (e.g., likes and responses received). Feedback features like the number of received likes and comments were also described as a measure for bonding social capital in previous work (Bohn et al., 2014). Based on these observations, we selected the most suitable feedback features as ground truth approximations in the following experiments. To allow further insights, two different ground truths approximations were investigated: a direct measure (e.g., number of likes) and a normalized version of it (e.g., number of likes divided by number of posts written). The motivation to use the normalized ground truth approximations is previous work that used ratios to approximate influence, which is related to CSC (Bentwood, 2008) (Anger and Kittl, 2011).

12.5 CSC ANALYSIS ON FACEBOOK

12.5.1 *Description of the Data Set*

The data set was crawled in the context of the bachelor's thesis of Monika Ullrich, which was supervised by Jan Hauffa and Georg Groh. The data set contains data of 11,629 users. For each user the data set includes personal information and interaction data in the form of posts, comments, friendships, and likes. These entity relationships are visualized in figure 12. More details can be found in (Ullrich, 2014).

12.5.2 *Selection of a Ground Truth Approximation*

Two different values were used as ground truth approximation for a person's CSCW in the investigation of the Facebook data set. The first value was $B_1 =$ Number of likes a person received, which was the most important value for the prediction of CSCW in the social capital experiment (see table 6). The Facebook data set includes user data that was created over several years. Some users have written several thousand posts in this period. The total number of likes some of these people received is high, even though the average number of likes per post is low. In order to treat people who registered later in a fair way, we used a normalized version of the count of likes as second ground truth approximation $B_2 = \frac{\text{Number of likes}}{\text{Number of posts} \cdot \text{Number of friends}}$. The number

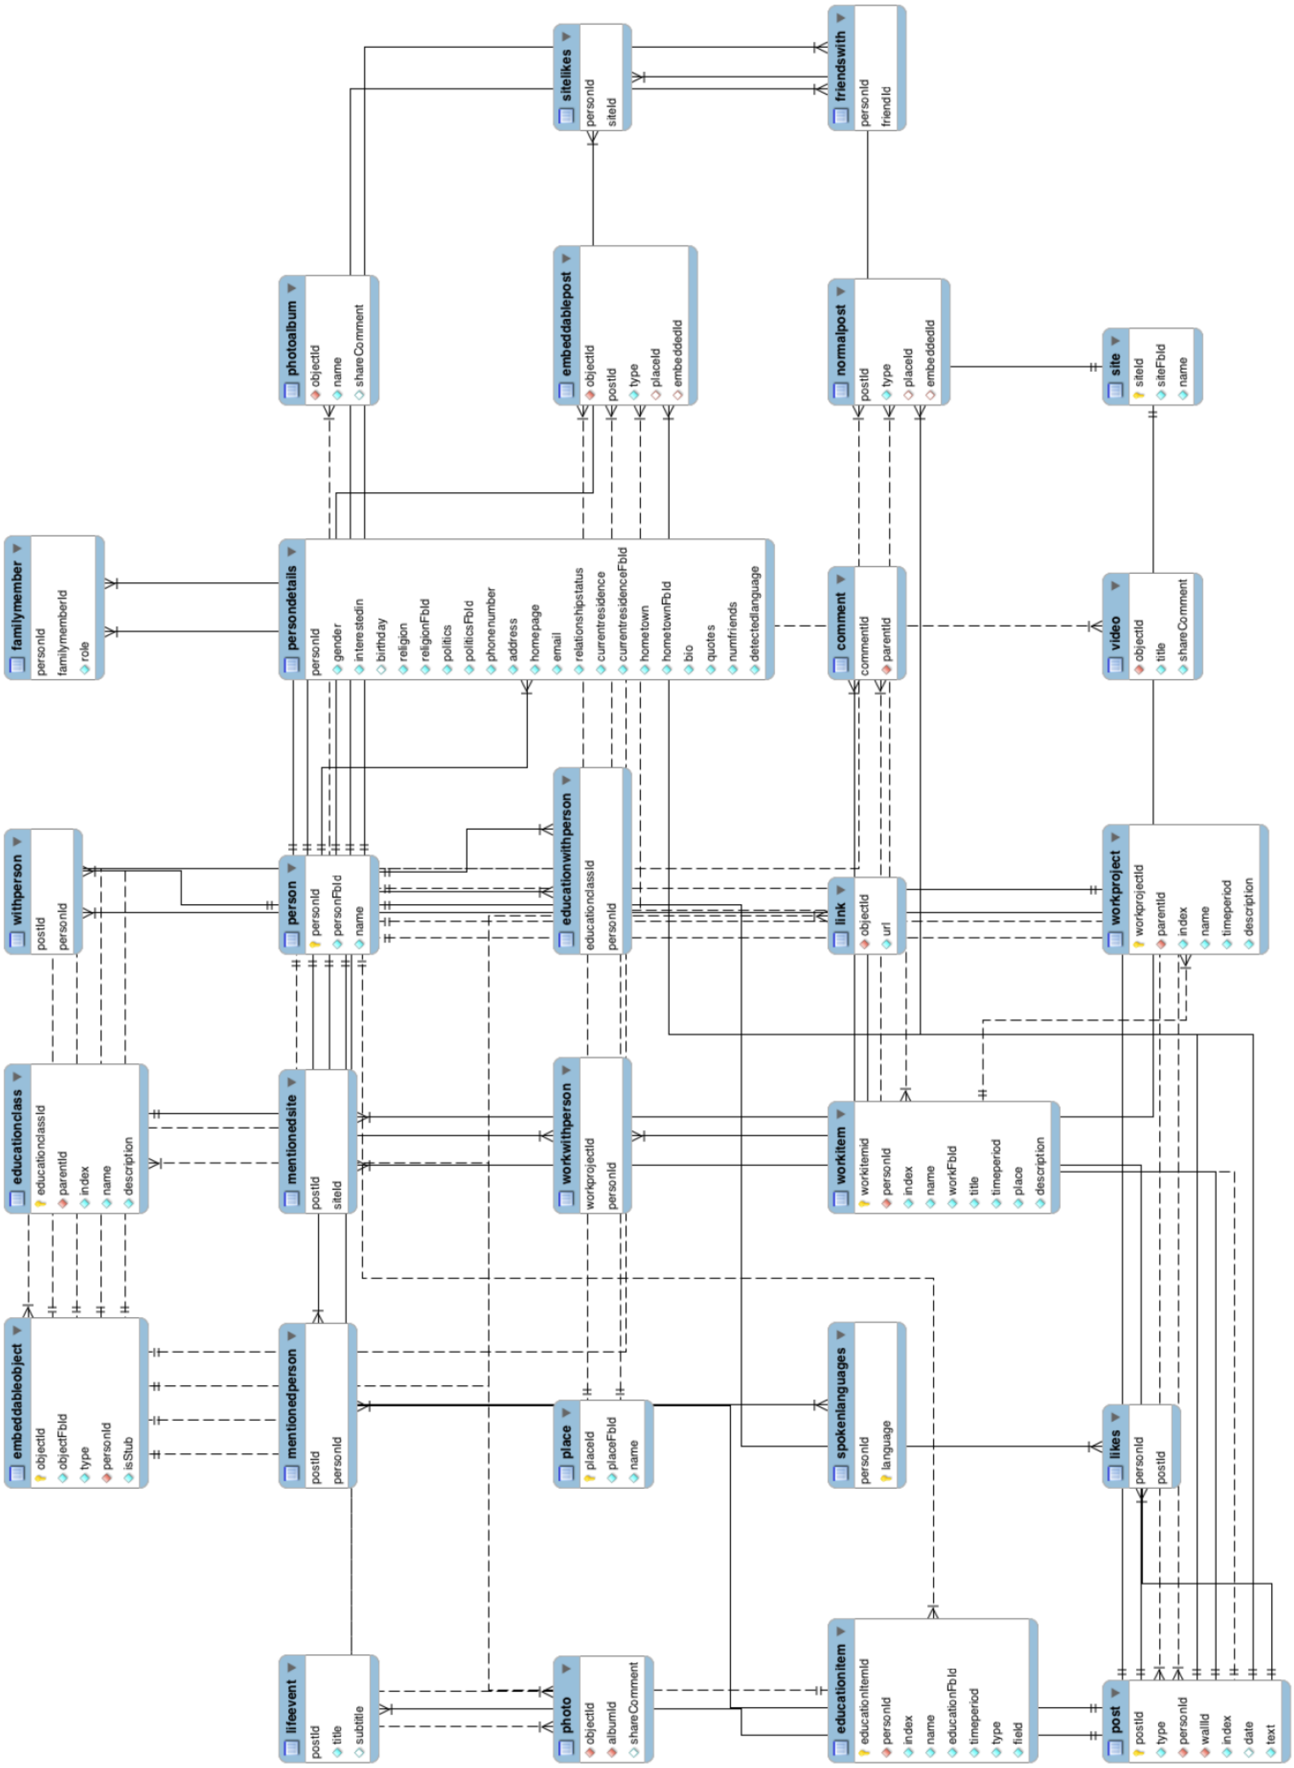

Figure 12: Entity relationship model of Facebook data set (Varshney, 2017)

of posts and friends in the denominator accounts for the activity of the user and the audience and thereby creates a ground truth approximation that allows a comparison of users solely based on the positive feedback they receive per post and friend.

12.5.3 *CSC Analysis*

The first step of the analysis was the identification and definition of features that could be used as input for the supervised learning algorithms. We selected and calculated as many features as pragmatically accessible from the respective data sources. For the selection privacy, availability, and a reasonable computability were taken into account. In section 12.5.4, the most and least important features are discussed, which allows a prioritization for future experiments. The features can be divided into four categories.

The first category are **user information features** that can be found on a user's profile page:

- **Number of friends.**
- **Highest level of education.** A 1 is assigned if the highest level of education is "highschool", a 2 if the education is a bachelor's degree or similar, 3 if it's a master's degree or equivalent, and a 4 if the person holds a PhD. If no education is stated and no education level can be inferred from the education institute, the value 0 is assigned.
- **Education information provided.** Is set to 0 when the education entry is empty and 1 otherwise.
- **Number of languages.** The number of languages that is listed by the user.
- **Language information provided.** Is set to 1 when the user specified at least one language and 0 otherwise.

The second set of features are **activity features** that reflect a person's activity within the network:

- **Number of posts.** Overall number of posts written by the user. The following features further describe the different types of posts that were used for the analysis.
- **Number of life event posts.** Number of life events (e.g., wedding) a user posted about.
- **Number of small posts.** Small posts are posts that are automatically grouped by Facebook on the wall. These are essentially birthday wishes.
- **Number of normal posts.** Number of all posts that are neither life events nor small posts.
- **Media post ratio.** Ratio of posts containing multimedia data (e.g., pictures or videos).

- **Average normal post length.** Average number of characters in normal posts.
- **Friends per post.** Ratio of the number of a user's friends to the number of written posts.
- **Number of comments.** Comments are answers on own or other people's posts.
- **Average length of comments.** Average number of characters in comments.
- **Number of users liked.** Number of distinct users whose posts or comments were "liked" by the user.
- **Number of likes.** Number of times the user "liked" posts or comments.
- **Number of users commented on.** Number of distinct users on whose posts the user commented on.
- **Number of comments on own posts.**
- **Number of comments on own life events.**

The third category are **feedback features** that reflect the responses and feedback a user creates in the network:

- **Number of likes.** The number of times a person's comments or posts are "liked" by others.
- **Likes per contribution.** Number of likes divided by number of comments and posts.
- **Unique mentions.** Number of people who mentioned the user in their posts and comments.
- **Mentions per friend.** Ratio of times the user is mentioned in other people's posts and comments to the number of friends.
- **Number of comments received.** Number of comments received on posts.
- **Users who commented.** Number of distinct users who commented on posts by the user.

The fourth category of features are centrality measures that reflect the network structure and the position of a user within the network. For Facebook we calculated 20 **centrality measures** that were extracted from five different graphs. For each graph the closeness, betweenness, eigenvector, and PageRank centrality (see section 4.1.2) were calculated for each user. The different graphs are:

- **Like graph unweighted.** The nodes are the users and an edge exists between all users who exchanged a like.
- **Like graph weighted.** Similar to the previous graph but with a weight at each edge that reflects the number of likes that were exchanged.
- **Friendship graph.** The nodes are users connected by edges that represent the friendship relationships between them.

Algorithm	R^2	Mean absolute error	Pearson r	Spearman ρ
Linear regression	0.870	0.046	0.933	0.936
Linear regression with regularization	0.847	0.052	0.921	0.923
Decision tree	0.843	0.051	0.918	0.912
Random forest	0.892	0.041	0.945	0.946
Neural network	0.876	0.045	0.937	0.941

Table 7: Performance of the different algorithms when predicting B_1 as ground truth approximation on the Facebook data set, compared according to four performance measures.

- **Comment graph unweighted.** An edge exists between all users who exchanged at least one comment.
- **Comment graph weighted.** Between all users who exchanged a comment, an edge exists that is weighted by the number of exchanged comments.

As the range of the values of features is large, the next step was a feature transformation. This was achieved by taking the logarithm and scaling the feature value to the range $[0, 1]$. This can be described with the following two equations.

$$x'_i := \log(x_i + 1) \tag{44}$$

$$x''_i := \frac{x'_i - x'_{min}}{x'_i - x'_{max}} \tag{45}$$

where $x = (x_1, ..., x_n)$ are the values for a feature for all samples $i \in \{1, ..., n\}$ and $x_{max} = \max_{i \in \{1,...,n\}} x_i$ and $x_{min} = \min_{i \in \{1,...,n\}} x_i$ are the respective maximum and minimum values. The predictions with and without transformation were compared and the evaluation parameters of the transformed features exhibited a significant improvement.

With the help of the transformed features and the ground truth approximations, four different supervised learning algorithms were trained on a training set that contained 80 percent of the data. Both ground truth approximations were used separately and the respective features – including the centrality measures calculated from them – were excluded from the input. The algorithms were then evaluated along different criteria on a test set that contained the remaining 20 percent of the data.

12.5.4 *Discussion of Results*

The results for the two different ground truth approximations are listed in tables 7 and 8.

The average R^2 value when using B_1 = Number of likes as ground truth approximation is 0.87. That means about 87 percent of the variance can be explained by

Algorithm	R^2	Mean absolute error	Pearson r	Spearman ρ
Linear regression	0.717	0.059	0.847	0.832
Linear regression with regularization	0.633	0.067	0.807	0.788
Decision tree	0.554	0.073	0.748	0.720
Random forest	0.724	0.058	0.853	0.839
Neural network	0.771	0.053	0.880	0.867

Table 8: Performance of the different algorithms when predicting B_2 as ground truth approximation on the Facebook data set, compared according to four performance measures.

the model. The average mean error value of 0.05 for B_1 as ground truth is relatively low and indicates that a prediction of B_1 is possible. The correlation coefficients underline these findings. The average Pearson r is 0.93 and the average Spearman ρ is also 0.93. Both values indicate a high correlation that is statistically significant at the <0.001 level. The best results were achieved with the random forest regressor. Its predictions had the highest R^2 value (0.89), the lowest mean error (0.04) and the largest correlation coefficients ($r = 0.95$ and $\rho = 0.95$).

Using $B_2 = \frac{\text{Number of likes}}{\text{Number of posts}\cdot\text{Number of friends}}$ as a ground truth, results in slightly worse results. The R^2 value decreases slightly, as this model is, on average, only able to explain 68 percent of the variance. The average mean error has a value of 0.06. The average correlation coefficients are slightly lower ($r = 0.83$ and $\rho = 0.81$) and also statistically significant at the <0.001 level. The best results were achieved with the neural network.

The differences between the two ground truth alternatives were also observed in the following three experiments, and are discussed consolidated in section 12.9.

Random forest and neural network were also the algorithms that performed best on the social capital experiment (compare chapter 12).

The most important features for the prediction of B_1 were the number of comments received (relative importance of 0.48), the media post ratio (0.08), and the number of normal posts (0.07). The features with the least influence on the prediction (relative importance of less than 0.01) are the education of a user, whether language information is provided and the number of languages, the unique mentions, and the number of mentions per friend. Similar to the social capital experiment, the most important features are those that reflect the engagement with the community (e.g., number of comments received), and the number of contributions to the network (e.g., number of posts). User information (e.g., languages, education) were less important.

For the prediction of B_2, the same features were in the top three, but this time they all had about the same importance. The media post ratio was the most important feature (relative importance of 0.22), followed by the number of normal posts (0.21), and the number of comments received (0.20). This time there were three features with limited importance for the model. All three features are again from the feature set user information: Education information provided, number of languages, and language information is provided.

For both predictions there are no centrality features among the most and least important features.

The conclusions regarding the overarching research questions (section 1.2) of the CSC system are discussed consolidated with the results of the other three experiments in section 12.9.

12.6 CSC ANALYSIS ON TWITTER

12.6.1 *Description of the Data Set*

The data set was crawled in the context of the master's thesis of Florian Hartl (Hartl, 2013). Twitter has a public API[2] that allows for a more comfortable crawling than on Facebook or Quora. Another aspect of relevance for our investigations is that a large percentage of communication is public and not restricted by privacy settings, which makes Twitter an ideal candidate for social capital analysis. The data set was crawled with a procedure based on Breadth-First-Search (Granovetter, 1976), which started with a random English-speaking and active Twitter user. The complete data set contains 358,342 users and over 220 million tweets. The data set was preprocessed to discard users who were inactive, just registered, and did not post in English. The data set used for the analysis included 25,000 users and 16,563,759 tweets.

12.6.2 *Selection of a Ground Truth Approximation*

The feature that was of the highest importance for the ground truth prediction in the social capital experiment was the number of likes received on comments (see table 6). The number of likes is a measure for how much the other users agree with the statement. On Twitter, there is a measure that is an even larger sign for agreement and appreciation of a user's statement: "retweeting" the comment of a user. Consequently, this feature was used as first ground truth alternative: $B_1 = \text{Number of retweets}$. The number of followers on Twitter varies even more than the friends' count on Facebook. To account for this variance, a second ground truth alternative was used that normalizes the number of retweets: $B_2 = \frac{\text{Number of retweets}}{\text{Number of followers} \cdot \text{Number of tweets}}$.

12.6.3 *CSC Analysis*

The analysis of the Twitter data set was conducted following the same steps of the Facebook analysis. At first we identified and computed features from the data set. The features were divided into four categories, the first being **user information features**:

- **Number of followers.** The number of the user's followers.
- **Days since sign-up.**

The second set of selected features reflects a person's **activity** in the network:

2 https://developer.twitter.com/en/docs (retrieved 2018-05-15)

- **Number of tweets.** Number of tweets posted by the user.
- **Average number of tweets per day.**
- **Average length of tweets.** Average number of characters in tweets.
- **Number of mentions in tweets.** The number of times an "@" appears in all tweets.
- **Number of tweets with mentions.** The number of tweets that have at least one "@" in them.
- **Hashtags in tweets.** The number of times an "#" appears in all tweets.
- **Number of tweets with hashtags.** The number of tweets that have at least one "#" in them.
- **Number of retweets.** Number of times the user retweeted other people's tweets.
- **Number of users retweeted.** Number of distinct users whose tweets were retweeted at least once by the user.
- **Number of replies sent.**
- **Number of users replied to.**

The third set of features reflect the **responses** and feedback a user receives in the network:

- **Number of followees.** Number of people who follow the user.
- **Received retweets.** Number of times the user's tweets were retweeted.
- **Number of users who retweeted.** Number of distinct users who retweeted.
- **Number of received replies.**
- **Number of users who replied.**

With the help of different graphs that were extracted from the Twitter data set, additional features were created. The **centrality measures** used were closeness, betweenness, eigenvector, and PageRank centrality. The following graphs were used for the extraction of these measures.

- **Follower graph.** Nodes are the users and the edges are follower relationships.
- **Retweet graph unweighted.** Nodes are the users and an edge exists between them if at least one retweet happened.
- **Retweet graph weighted.** The edges between the users are weighted by the number of retweets exchanged between them.
- **Reply graph unweighted.** Nodes are the users and an edge exists between them if at least one user replied to the other.

Algorithm	R^2	Mean absolute error	Pearson r	Spearman ρ
Linear regression	0.782	0.069	0.884	0.843
Linear regression with regularization	0.686	0.086	0.857	0.822
Decision tree	0.765	0.067	0.875	0.828
Random forest	0.822	0.058	0.907	0.856
Neural network	0.816	0.060	0.904	0.851

Table 9: Performance of the different algorithms when predicting B_1 as ground truth approximation on the Twitter data set, compared according to four performance measures.

Algorithm	R^2	Mean absolute error	Pearson r	Spearman ρ
Linear regression	0.115	0.007	0.348	0.484
Linear regression with regularization	0.116	0.007	0.345	0.483
Decision tree	0.179	0.005	0.446	0.451
Random forest	0.251	0.004	0.527	0.600
Neural network	0.135	0.006	0.373	0.485

Table 10: Performance of the different algorithms when predicting B_2 as ground truth approximation on the Twitter data set, compared according to four performance measures.

- **Reply graph weighted.** The edges between the users are weighted by the number of replies exchanged between them.

All of these features were transformed following equations 44 and 45. For the analysis, the features that represent the selected ground truth alternative directly or indirectly were disregarded.

12.6.4 *Discussion of Results*

Tables 9 and 10 list the results of the prediction of the two different ground truth approximations.

The prediction of $B_1 =$ Number of retweets, which is visualized in table 9, works quite well — the average R^2 value is 0.77, which indicates that almost 80 percent of the variance can be explained by the model. The mean average error is 0.07 and thereby higher than for the corresponding ground truth approximation in the Facebook data set but still relatively small. The Pearson and Spearman correlation coefficients are high ($r = 0.89$ and $\rho = 0.84$ with a p value of <0.001). All these results indicate that B_1 can be predicted well with the help of the described features. The best algorithm for the prediction was the random forest regressor — it had the highest R^2 value (0.82), the lowest mean average error (0.06) and the highest correlation coefficients ($r = 0.91$ and $\rho = 0.86$).

The prediction of $B_2 = \frac{\text{Number of retweets}}{\text{Number of followers}\cdot\text{Number of tweets}}$ is less clear, mainly because the coefficient of determination R^2 is much lower. On average the model can only explain 16 percent of the variance, which is a small value for such a large data set. The absolute mean error is 0.01 and thereby only about a seventh of the mean absolute error of the B_1 prediction. However, the mean value of the B_1 ground truth approximation is about 40 times larger than the average B_2, after scaling and normalization. The relative error of the B_2 prediction is consequently larger than the relative error of the B_1 prediction. The correlation coefficients indicate that there is a mediocre correlation between the predicted values and the ground truth approximations ($r = 0.41$ and $\rho = 0.50$). This is significantly worse than for the prediction of B_1. The difference between the two ground truth variables with regard to the predication performance is discussed in section 12.9. The best algorithm for the predication of B_2 is the same as for the prediction of B_1: the random forest regressor.

If one disregards the centrality features, the most important features for the prediction of B_1 are the number of users who replied (relative importance of 0.28), the number of followers (0.19), and the average length of the tweets (0.8). These are all activity and feedback features that reflect how active a user is and from how many users he or she is supported. This is supported by previous work from Hofer and Aubert who found that the number of followers is an important indication for a person's social capital (Hofer and Aubert, 2013). The features with low relevance are the number of hashtags in tweets, the number of users replied to, and the number of tweets with hashtags. The low importance of the features with hashtags may be explained by the fact that a low percentage of tweets on Twitter contain hashtags (Kywe et al., 2012). The fact that the number of people the user replied to has little relevance for the prediction of the number of retweets, indicates that retweets cannot be forced by activity alone. By including the centrality features, the picture changes slightly. The number of users who replied is still the most important feature (0.16), followed by the PageRank centrality on the unweighted follower graph (0.12), and the number of followers (0.08). All centrality measures calculated on the unweighted reply graph are of little importance for the prediction of B_1.

The most important features for the prediction of the normalized ground truth approximation $B_2 = \frac{\text{Number of retweets}}{\text{Number of followers}\cdot\text{Number of tweets}}$ were the average number of tweets per day (relative importance of 0.30), the number of days since sign-up (0.22), and the number of followees of a user (0.10). It is possible to reconstruct the number of tweets written by a user from the first two features. The number of tweets was used for the normalization of the ground truth value B_2. This may explain their relative importance. If these two features are disregarded, the features of the highest relative importance are the number of followees, the average length of the tweets, and the number of tweets with hashtags. The number of followees is a feedback feature and the number of tweets per day, as well as the features regarding tweet length and number of hashtags are activity features illustrating the activity of a user. The least important features are all centrality features from the reply graph. Disregarding the centrality measures leaves the number of received replies, the number of users who were retweeted by the user, and the number of people the user replied to as least important.

The conclusions regarding the overarching research questions are discussed consolidated with the results of the other three experiments in section 12.9.

12.7 CSC ANALYSIS IN SCIENTOMETRICS

12.7.1 *Description of the Data Set*

Citation networks are different from the other networks investigated in this chapter, not least because they can only reference backwards and do not directly reflect communication between people. However, the publications that reference each other reflect the author's expertise, a vital component of their CSC. ArnetMiner[3] is a platform with the goal of "extracting and mining academic social networks" (Tang et al., 2008), which we can leverage for CSC analysis. The profiles of scientists on this platform were created with information from several digital libraries, like the ACM Digital Library[4], the DBLP Computer Science Bibliography[5], or CiteSeer[6], as well as information from the personal websites of the scientists. ArnetMiner contains mainly authors in the field of computer science and mathematics and provides data sets for research purposes free of charge. We chose a large academic social network[7] with more than 1.7 million authors and over 2.0 million papers. The citation network has over 8.0 million edges between papers and the co-author network over 4.2 million collaboration relationships. This data set was preprocessed in order to exclude authors that could not be used for our analysis. We excluded all authors whose information were not complete, e.g., because they did not have at least one valid abstract. The valid abstract was important because we used topic modeling to identify topical interests of the authors based on these abstracts. The preprocessed data set included 99,178 authors and 559,717 papers, and 2,331,154 citations.

12.7.2 *Selection of a Ground Truth Approximation*

A universal token for appreciation in the world of science is a citation. It reflects that one read and understood the publication of a fellow scientist and values their work so highly that one builds a part of the own research on the article. The count of citations consequently reflects how important the contributions of an author were to the network. Therefore, we select $B_1 = \text{Number of citations}$ as first ground truth approximation. To adjust for the fact that scientists with more publications usually also accumulate more citations, two additional ground truth approximations were investigated: the h-index (see section 10.4) and a normalized version of citations $B_2 = \frac{\text{Number of citations}}{\text{Number of papers}}$. Both ground truth approximations set a person's number of citations in relation to their published papers. The h-index additionally accounts for

3 http://www.arnetminer.cn (retrieved 2018-05-15)
4 https://dl.acm.org/ (retrieved 2018-05-15)
5 https://dblp.org/ (retrieved 2018-05-15)
6 http://citeseerx.ist.psu.edu/ (retrieved 2018-05-15)
7 https://www.aminer.cn/aminernetwork (retrieved 2017-06-10)

people who only have one highly cited article and several unrecognized ones and ranks them lower than people who achieved medium citation numbers on several papers. While both approaches produce values that certainly correlate with a person's CSC, it is difficult to decide which one is the better approximation. We investigated both options and obtained similar results. For symmetry reasons with the other experiments in this chapter, we will only discuss the results of the ground truth alternative $B_2 = \frac{\text{Number of citations}}{\text{Number of papers}}$ in the following.

12.7.3 *CSC Analysis*

The features in the ArminMiner database are different from the two previous experiments. That is due to the nature of the network that reflects the scientific exchange with a much lower exchange frequency of contributions. Nevertheless, one can divide the features in the same four categories.

Features that describe **user information** are:

- **Number of institutions.** Number of universities and other institutions that are listed in the author's profile.
- **Number of institutions in top 500.** Based on input from the Academic Ranking of World Universities[8] (also known as Shanghai ranking), the National Taiwan University (NTU) Ranking[9], and the Times Higher Education (THE) Ranking[10] we created a list of the top 500 universities in the world. This feature counts the number of institutions an author is associated with that are in this list.
- **Shanghai rank.** The rank of the author's institution in the Shanghai ranking. If there is more than one institution, the lowest (best) score is selected.
- **NTU rank.** The rank of the author's institution in the NTU ranking. If there is more than one institution, the lowest (best) score is selected.
- **THE rank.** The rank of the author's institution in the THE ranking. If there is more than one institution, the lowest (best) score is selected.
- **Shanghai score.** The score the author's institution received in the Shanghai ranking (maximum of 100 points). If there is more than one institution, the highest (best) score is selected.
- **NTU score.** The score the author's institution received in the NTU ranking. If there is more than one institution, the highest (best) score is selected.
- **THE score.** The score the author's institution received in the THE ranking. If there is more than one institution, the highest (best) score is selected.
- **Topic similarity to full data set.** The topic similarity of the author's topic distribution θ_i and the topic distribution of the entire corpus θ_C. This can be interpreted as a measure for how much the author publishes in topics that are of interest to the majority in the network.

8 http://www.shanghairanking.com (retrieved 2017-11-14)
9 http://nturanking.lis.ntu.edu.tw/Default.aspx (retrieved 2017-11-14)
10 30https://www.timeshighereducation.com/world-university-rankings (retrieved 2017-11-14)

- **Topic similarity to uniform topic distribution.** The topic similarity of the author's topic distribution θ_i and the uniform topic distribution θ_U. This is a measure of diversity: the smaller the distance to the uniform distribution, the more topics the author is interested in and the more diverse the author is.
- **Topic similarity of papers.** For each of the author's papers the topic similarity to all other papers is computed and then averaged for every author. This is also a measure for diversity.
- **Number of topics greater than corpus.** We compare the author's topic distribution θ_i and the topic distribution of the whole corpus θ_C and count the number of times the value of the author is greater than the value of the entire corpus. This feature reflects how many topics the author is involved with.
- **Number of topics greater than uniform.** We compare the author's topic distribution θ_i and the uniform topic distribution θ_U and count the number of times the value of the author is greater than the value of the uniform distribution. This is also a measure for how many different topics the author is involved with.

Activity features are:

- **Number of publications.** Number of publications of which the person was an author.
- **Number of first positions.** The number of times the person was the first author of a paper.
- **Number of second positions.** The number of times the person was the second author of a paper.
- **Number of third positions.** The number of times the person was the third author of a paper.
- **Number of higher positions.** The number of times the person was mentioned at the fourth or a higher position of a paper.
- **Years active.** Number of years since the author's first publication.
- **Years between first and last.** Number of years between the author's first and last publication.
- **Average number of publications per year.**
- **Average title length.** Average number of characters in the titles of the author's publications.
- **Average abstract length.** Average number of characters in the abstracts of the author's publications.
- **Number of authors referenced.** Number of distinct authors whose publications were cited by the author.

- **Number of references.** Total number of citations by the author.

The following **feedback features** were investigated in the ArminMiner database:

- **Number of citations.** Number of times the author was cited by other authors.
- **Number of collaborating authors.** The number of distinct authors the author collaborated with.
- **Number of collaborations.** The number of times the author collaborated with others.

Similar to the previous experiments, closeness, betweenness, eigenvector, and PageRank **centrality** were used on different graphs extracted from the data set. Four graphs were used for the analysis:

- **Collaboration graph unweighted.** The nodes are the authors and an edge exists between all authors who collaborated on at least one publication.
- **Collaboration graph weighted.** Similar to the previous graph but with a weight at each edge that reflects the number of times the authors collaborated.
- **Citation graph unweighted.** The nodes are the authors and an edge exists between all authors who were cited at least once by one another.
- **Citation graph weighted.** Similar to the previous graph but with a weight at each edge that reflects the number of times the authors were cited by one another.

Similar to the previous experiments, all of these features were transformed following equations 44 and 45. Only the features that do not include the respective ground truth alternative were used for the analysis.

12.7.4 *Discussion of Results*

We compared the prediction results of the different algorithms in tables 11 and 12, respectively for B_1 and B_2.

The observations that we can make about the results of the prediction are in line with the two previous sections about Facebook and Twitter.

The prediction of the quantitative ground truth approximation $B_1 =$ Number of citations works relatively well. The average R^2 value of 0.79 indicates that the model can explain most of the variance. The mean error is 0.07 and thereby in the same range as the predictions on Facebook and Twitter. The high average correlation values ($r = 0.89$ and $\rho = 0.86$) underline the assumption that the prediction of the number of citations, which is likely correlated with an author's CSC, works well with the identified features and the algorithms. The best algorithm for the prediction is again the random forest regressor.

The prediction of the normalized version of the ground truth approximation $B_2 = \frac{\text{Number of citations}}{\text{Number of papers}}$ performs less well. The coefficient of determination has an average value of $R^2 = 0.43$, i.e. less than half of the variance can be explained by the model.

Algorithm	$\mathbf{R^2}$	**Mean absolute error**	**Pearson** r	**Spearman** ρ
Linear regression	0.781	0.073	0.884	0.853
Linear regression with regularization	0.752	0.080	0.876	0.848
Decision tree	0.785	0.072	0.886	0.859
Random forest	0.827	0.063	0.909	0.889
Neural network	0.804	0.069	0.898	0.870

Table 11: Performance of the different algorithms when predicting B_1 as ground truth approximation on the ArnetMiner data set, compared according to four performance measures.

Algorithm	$\mathbf{R^2}$	**Mean absolute error**	**Pearson** r	**Spearman** ρ
Linear regression	0.386	0.080	0.621	0.626
Linear regression with regularization	0.359	0.082	0.602	0.615
Decision tree	0.417	0.076	0.647	0.639
Random forest	0.528	0.067	0.727	0.716
Neural network	0.446	0.073	0.685	0.668

Table 12: Performance of the different algorithms when predicting B_2 as ground truth approximation on the ArnetMiner data set, compared according to four performance measures.

The average mean error is slightly larger than for the B_1 prediction; 0.08. The correlation coefficients also indicate that the prediction does not work as well ($r = 0.66$ and $\rho = 0.65$).

The reasons for this difference to the B_1 prediction are discussed in section 12.9.

The best algorithm for the prediction is the random forest regressor.

The most important feature for the prediction of B_1 = Number of citations, is the number of publications (relative importance of 0.62). The number of references used by the author, is the second most important feature (0.08), followed by the number of years since the first publication (0.07). The number of publications and the number of active years are correlated and signs for the overall activity of the author. The importance of the number of references may be an indication for the existence of reciprocity — that people reference others who previously cited them. There are no centrality measures among the most important features. The least important features without centrality measures are the number of topics greater than uniform, which is a measure that reflects how many different topics an author is interested in, the number of institutions that are in the top 500 list, and the number of institutions affiliated with the author. Including the centrality features reveals that the betweenness and eigenvector centralities of the collaboration graph are of little importance.

For the prediction of $B_2 = \frac{\text{Number of citations}}{\text{Number of papers}}$, we first investigate the most and least important features without the centrality measures. The number of years the author is active is the most important feature (relative importance of 0.14), closely followed by the number of references (0.14). The number of years between the first and last publication is also important (0.09). The duration features are correlated with each other and with the number of publications, which is in turn the most important feature for the prediction of the number of citations. The least important features are the number of times the person was the second author of a publication, and again the number of institutions that are in the top 500 list, and the number of institutions. By including the centrality measures, the picture changes. The most important feature is the closeness centrality of the unweighted collaboration graph (0.18), which indicates that being close to other scientists, something that may be achieved by networking, is of high importance for a high B_2 prediction. This is followed by the number of active years (0.11) and the years between the first and last publication (0.07). The least important features are still the number of institutions and the number of institutions in the top 500. The betweenness centrality of the weighted collaboration graph is of the third least importance. None of the five features that contained topic information were of significant importance for the prediction of both ground truth approximations. This indicates that the topical field in which an author publishes, and their personal writing style as assessed by the topic model, are not directly relevant for the number of citations they receive.

Conclusions about the overarching research questions stated in section 1.2 are discussed consolidated with the results of the other three experiments in section 12.9.

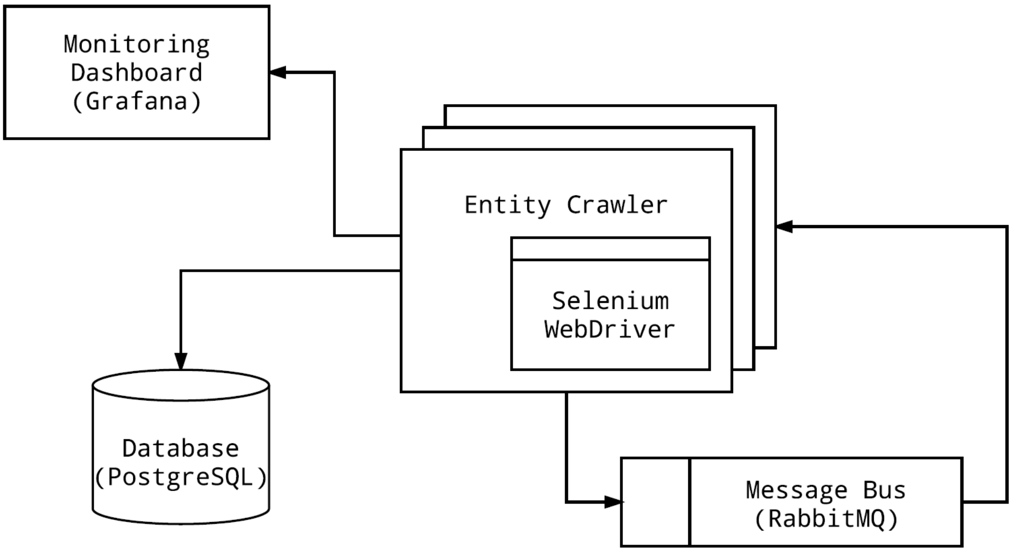

Figure 13: Diagram of the Quora crawler architecture (Chernenko, 2017)

12.8 CSC ANALYSIS ON QUORA

Quora[11] is a question and answer portal with no specific topic focus. Similar to many threaded discussion boards, users can ask and answer questions, converse with each other (usually about the topics of the question), and evaluate answers by up-voting and down-voting. Quora also includes elements of a social networking platform and allows users to follow each other.

12.8.1 *Description of the Data Set*

The data set was crawled in the context of Valeriia Chernenko's master's thesis in 2017 (Chernenko, 2017). This thesis was supervised by Sebastian Schams and Georg Groh. In the following, we briefly summarize the crawling process. The crawler consists of a database, several entity crawlers, and a message bus, as visualized in figure 13.

With the help of a Selenium WebDriver[12], each entity crawler can load and navigate web pages. There are four different types of entity crawlers:

- *Question crawler:* Saves the title and all topic tags of a question and all answers posted to it. It sends requests to the user crawlers to crawl basic statistics of all users it encounters. For every answer the text and media content as well as the name and credentials of the user answering the question are saved. It also saves a time stamp and the number of views and up-votes of the answer.

- *User information crawler:* Saves the information that is provided on each user's profile page. This information includes the name, a description, and basic statistics, like the number of answered questions and followers. It can send requests to the user profile crawler to collect more information.

11 https://de.quora.com/ (retrieved 2018-05-15)
12 1http://www.seleniumhq.org/projects/webdriver/ (retrieved 2018-05-15)

- *User profile crawler:* The complete user profile includes a list of all questions and answers of the user, as well as the topics they are interested in. This information is collected by this entity crawler, together with a list of the user's followers and followees. A request is sent to the user information crawler to collect basic information about these people. Additionally, a request is sent to the topic crawler to collect further information about the topics of interest.

- *Topic crawler:* Collects the number of questions that are tagged with this topic and the number of user's who follow this topic.

With the help of the message bus, the entity crawlers can create requests for each other and put them in the respective queues. This architecture enables a fast and effective crawling process, as it allows to simply create more entity crawlers of a specific type to prevent bottlenecks. The Quora crawler ran over a period of three weeks on a machine with 12 Intel(R) Core(TM) i7-3930K CPU @ 3.20GHz processors and 58GB RAM. It crawled over 53 thousand questions with 817 thousand answers, 213 thousand (partial) user profiles, and almost 122 thousand topics. After preprocessing, e.g., reducing users whose profiles were only partially crawled, we obtained a data set of 3,069 users, who posted 143 thousand questions and answered 36 thousand questions.

12.8.2 *Selection of a Ground Truth Approximation*

Useful contributions in the form of answers or interesting questions are recognized by the Quora community with up-votes. Up-votes are similar to likes on other networks and, therefore, the first choice for a ground truth alternative is B_1 = Number of up-votes. Even though follower relationships are possible on Quora, most people select questions based on their topic interests. The number of views a question or answer has, therefore, is likely the best metric to normalize the number of up-votes: $B_2 = \frac{\text{Number of up-votes}}{\text{Number of views}}$.

12.8.3 *CSC Analysis*

As was the case in the previous three experiments, the features that can be identified in the data set crawled from Quora can be categorized in four different sets.

User features are:

- **Number of topics.** The number of topics the user follows.

- **Followers and followees ratio.** The ratio of the number of people who follow the user to the number of people the user follows.

- **Top score ratio.** The number of questions where the user provided the answer with the highest score divided by the number of all questions the user has answered.

- **Objective questions ratio.** With the help of a subjectivity classifier, all answers were classified in two categories: objective (e.g., "What year did France win the football world championship?") and subjective (e.g., "Who is your favorite musician?") questions. This ratio is the ratio of objective questions a user answered to the total number of questions the user answered.

- **Subjective questions ratio.** The ratio of subjective questions a user answered to the total number of questions the user answered.

Activity features are:

- **Number of answers.** The number of answers that were written by the user.
- **Number of questions.** The number of questions asked by the user.
- **Number of edits.** The number of edits the user made on answers or questions.
- **Number of followees.** The number of people the user is following.
- **Number of answers by others.** The total number of answers to questions that the user answered.
- **Number of users up-voted.** The number of distinct users that were up-voted by the user at least once.
- **Number of up-votes.** The total number of times the user up-voted answers.
- **Number of users whose questions were answered.** The number of distinct users whose questions were answered by the user at least once.
- **Average answer length.** Average number of words in an answer by the user.
- **Average answer SMOG index.** The SMOG index is a measure for the readability of written text (Mc Laughlin, 1969). It estimates the number of years of education required to understand the text. This feature is the average SMOG index of all of the user's answers.
- **Number of answers with a URL.** The number of answers with an URL, which may be an indication for a provided source.
- **Average subjectivity probability.** The average probability that a question answered by the user is subjective, which was determined by the subjectivity classifier.

Feedback features on Quora are:

- **Number of followers.** Number of people who follow the user.
- **Number of users who answered.** Number of distinct users who have answered questions by the user.
- **Number of answers received.** Total number of answers received on all questions by the user.
- **Number of views of answers.** The number of views the answers of the user have.

Algorithm	R^2	Mean absolute error	Pearson r	Spearman ρ
Linear regression	0.931	0.042	0.965	0.957
Linear regression with regularization	0.930	0.042	0.965	0.956
Decision tree	0.915	0.046	0.957	0.945
Random forest	0.932	0.041	0.965	0.956
Neural network	0.933	0.041	0.967	0.958

Table 13: Performance of the different algorithms when predicting B_1 as ground truth approximation on the Quora data set, compared according to four performance measures.

In order to leverage additional network information, five different graphs were extracted from the Quora data set. Closeness, betweenness, eigenvector, and PageRank **centrality features** were extracted from the following graphs for each user:

- **Up-vote graph unweighted.** The users (nodes) are connected if they exchanged at least one up-vote.
- **Up-vote graph weighted.** Like the previous graph but with edges weighted by the number of up-votes between the users.
- **Follower graph.** The users (nodes) are connected if they follow one another.
- **Answer graph unweighted.** The users (nodes) are connected if one user answered at least one question of the other.
- **Answer graph weighted.** Like the previous graph but with edges weighted by the number of answers between the users.

All of these features were transformed following equations 44 and 45 and the features used as ground truth approximation were disregarded.

12.8.4 *Discussion of Results*

The prediction results of the different algorithms are compared in tables 13 and 14 for B_1 and B_2.

The number of up-votes (B_1) a person receives can be predicted well with the help of the features listed in the previous section. This is important for CSC research, as an up-vote is a sign that an answer of a question was useful to another person in the community, which reflects the value-ad of a person to the network and thereby their CSCW. The average R^2 value for the B_1 prediction is 0.93, which is the second score of all of our experiments. The mean error is relatively low (0.04) and the average correlation coefficients ($r = 0.96$ and $\rho = 0.95$) are also the highest in all four experiments. A reason for this satisfying performance on the Quora data set may be the content focus and nature of Quora. On Facebook and Twitter, users may be compelled to "like"

Algorithm	R^2	Mean absolute error	Pearson r	Spearman ρ
Linear regression	0.224	0.080	0.480	0.479
Linear regression with regularization	0.225	0.079	0.474	0.477
Decision tree	0.084	0.084	0.295	0.356
Random forest	0.312	0.074	0.564	0.542
Neural network	0.269	0.077	0.522	0.509

Table 14: Performance of the different algorithms when predicting B_2 as ground truth approximation on the Quora data set, compared according to four performance measures.

other people's posts or comments, simply because they are befriended in real life. On Quora, users are usually more interested in correct answers for their questions rather than pleasing others with up-votes. Similar observations about the less explicit nature of relationships were made on the similar platform Slashdot (Gómez et al., 2008). If this hypothesis holds true, threaded discussion boards and Q&A portals may be ideal sources for CSC analysis.

The best algorithms for the prediction was the neural network, closely followed by the random forest regressor.

The prediction of normalized ground truth approximation $B_2 = \frac{\text{Number of up-votes}}{\text{Number of views}}$ was less successful. The average R^2 value indicates that only 22 percent of the variance can be explained by the model. The mean error (0.08) is almost twice as high as for the B_1 prediction. The average correlation coefficients are also worse ($r = 0.47$ and $\rho = 0.47$). Reasons for this discrepancy are discussed in the following section 12.9.

The best algorithm for the prediction of B_2 were the random forest regressor and the neural network.

In the following discussion of the importance of the features, we first disregard the centrality measures and then discuss what changes when the centrality measures are included.

Without centrality features, the most important feature for the prediction of the number of up-votes (B_1) was the number of views of a user's answers (relative importance of 0.47). The reason is likely that even helpful answers cannot be up-voted if they are not perceived enough. For the context of CSC this means that people who engage in topics that are of the most relevance for the whole community create the most added value. The second most important feature is the number of answers by others on the questions the user answered to (0.14). This feature is closely related to the previous one, as it reflects the importance of the topic the user engages in. The third most important feature is the number of followers of the user (0.07), which indicates that a user has provided answers that were relevant enough for others to decide to follow the user. The three least important features with almost no relative importance were the number of up-votes by the user, the number of users who answered questions by the user, and the total number of received answers. By including the centrality measures, there are still no centralities among the top three most important features. The number of views of a user's answers is still the most important (0.45), followed by the number of answers to questions answered by the user (0.09), and the number of edits a user made (0.08). The third most important feature changed compared to the case where the centrality measures were disregarded. This is likely because the number of followers that was previously the third most important feature is now included in a variety of centrality measures calculated on the follower graph, which reduces the relative importance. The PageRank, betweenness, and eigenvector centrality of the answer graph are the least important features.

For the prediction of $B_2 = \frac{\text{Number of up-votes}}{\text{Number of views}}$ the most and least important features, when disregarding centralities, display a large overlap to the prediction of B_1. The most important features are the user's top score ratio (relative importance of 0.12), the number of the user's followers (0.10), and the number of answers by others on the question the user answered to (0.10). The least important features are the number of answers with a URL, which contradicts the hypothesis that these answers are important because they indicate a linked source, the number of users who answered questions by the user, and the total number of received answers. The ranking of the top three does not change when including centrality measures. The least important features also include the PageRank and eigenvector centrality of the answer graph, when the centrality features are used for the prediction.

In the following section 12.9 the results of all four experiments are summarized and discussed with regard to the CSC system and the overarching research questions.

12.9 SUMMARY AND DISCUSSION OF RESULTS

In this chapter, we investigated four different data sources that are typical representatives of online social interaction and knowledge exchange. In the following the findings are discussed with regard to the research questions and CSC system as a whole.

In all four instances we used data sets that contained user information, activity features, feedback features, and centrality measures. These data sets allowed for an investigation of CSC for each individual user (first research question), which is in line with the findings of the experiment in chapter 11 and previous work (see chapter 10). In the case of scientometrics, we could leverage the ArnetMiner data set. In the other three instances, the data needed to be crawled from the platforms. The crawling process was time-consuming and took three weeks for the Quora data set alone. For a continuous assessment of the CSC as first pillar of the system, this is too slow. It underlines the importance to obtain permission for the data mining, as discussed in 9.2.3, in order to leverage the APIs of the platforms instead of random crawls of the publicly available data (overarching research question 2 from section 1.2). Once the algorithm is trained, online learning algorithms can be used for an continuous assessment of the users' CSCW.

The second research question asked for features that could be used for the analysis. A variety of features were identified and computed on all data sources and roughly grouped in four categories: user information features, activity features, feedback features, and centrality measures. Including the centrality measures, we identified 45 features on Facebook, 38 on Twitter, 44 for scientometrics, and 41 on Quora.

A detailed discussion of the relative importance of the different features, which is related to research question 6, is provided in the respective sections. In general, the most important features are feedback and activity features. This makes sense, as it directly reflects the value creation in the network (activity) and the positive feedback resulting from it. This is also what we have seen in previous work, e.g., regarding the identification of experts on Twitter (Hadgu and Jäschke, 2014).

The centrality measures did not have a large impact on the prediction. The R^2 value without all centrality measures was on average less than 2 percent smaller than with the centrality measures and the mean error only improved slightly when the centrality measures were included in the analysis. The effect of the centrality features on the correlation coefficients was also only a small improvement of about 1 percent.

There are two reasons that may explain the limited importance of the centrality measures. First of all the features that were used for the calculation of the centralities, e.g., the number of likes or friendship relationships, were also used for the analysis. In addition to these features the centralities only add little additional value. Secondly, it indicates that for the prediction of the respective ground truth alternatives, the graph structure does not provide much new information.

Especially for large networks, the calculation of the centrality features is computation-intensive. Considering their limited value-add for the prediction, it may be reasonable to omit centralities in future experiments. Language-based features, like the SMOG index of a user's contributions, topic similarities, or whether Quora users asked subjective or objective questions, also added little to the prediction quality.

The third research question was what value can be used to approximate the real CSCW. We investigated a quantitative and a normalized ground truth approximation for each data source. The choice of the feature used as quantitative ground truth approximation B_1 was motivated by previous work and the findings of the social capital experiment described in chapter 11, where a ground truth assessment in the form of direct peer feedback was available. The normalized ground truth approximation B_2 also used the quantitative feature which was normalized using appropriate metrics. This allowed to investigate the differences between the two approximations. Both B_1 and B_2 are likely less reliable than the peer assessment used in the experiment described in chapter 11. The reason is that these approximations only indirectly include the aspects of competence, trustworthiness, and social responsibility. The indirect inclusion of these characteristics is due to the selected features that incorporate the feedback of others. However, they still allow several observations about the first pillar of the CSC system.

On all data sets it was possible to predict the ground truth alternatives. The quality of the prediction varied significantly between B_1, the quantitative value that reflects feedback from others, and B_2, the normalized value. The unweighted average of R^2 for all B_1 (B_2) predictions is 0.84 (0.37), of the mean error 0.06 (0.06), of the Pearson r 0.92 (0.59), and of the Spearman ρ 0.90 (0.61). All values are better for the prediction of B_1. The seemingly similar value of the mean error of the prediction is misleading. Due to the normalization, the value of B_2 is up to an order of magnitude smaller than the average value of B_1, even after the transformation (logarithm and scaling). The relative error for B_1 is consequently significantly smaller. These differences regarding the quality of the predictions may be due to two factors.

The first reason is the reduction of features. For the normalization of the ground truth approximation, some additional features are required, e.g., the number of posts and the number of friends. These features are no longer used as input variables to prevent that the algorithm only uses this feature for the prediction. As these features are of high importance for the prediction of the ground truth approximation B_1 and important aspects of a person's online persona, a significant part of the input for the model is missing, which may be reflected by the lower R^2 values.

The second reason is the origin of the B_2 ground truth approximation. The initial assumption was to create a normalized value that allows to easily compare users with different interaction frequencies and different numbers of total contributions. This was achieved by dividing a feedback feature, like the number of likes, by the numbers of a user's friends or posts. The truth is, however, that the features in the denominator are often signs for a user's value-ad to the network and should, therefore, directly correlate with the CSC, as we have seen in table 6, and not in a reciprocal way. This indicates that the selection of B_1 as ground truth alternative is better than B_2, which is additionally supported by the performance differences.

Based on these observations, the conclusion for future investigations is to use as few features as possible as ground truth approximation.

The fourth research question was whether CSCW can be predicted based on the extracted features. This is similar to the overarching third research question (section 1.2). We demonstrated with the help of the two ground truth alternatives that it is generally possible to infer new features/characteristics from the four data sources. The best results were achieved on the platform Quora, even though it was not the largest sample. The average coefficient of determination R^2 on Quora was large (0.93) and the average mean error low (0.04) compared to the other experiments. This may be due to the nature of interactions on Quora, where the focus lies on the content of contributions rather than the person who wrote them. This indicates the importance of Q&A platforms like Quora and Stack Overflow for the assessment of CSC.

The best algorithm for the prediction on all data sets was in most cases the random forest regressor (research question five). On Facebook and Quora the prediction of the neural network was slightly better in some instances. Multiple linear regression with regularization and decision trees were usually the algorithms that performed the worst. This indicates that more complex algorithms are required in order to obtain the best results.

The reasoning is similar to the one of the experiment described in chapter 11. The low performance of multiple linear regression indicates that the relationship between the ground truth assessment and the respective features cannot be described by a linear model. The non-linear relationship can be learned by both random forest and neural network. The neural network can express the non-linearity directly due to its non-linear activation function. The random forest approximates the non-linearity in a piecewise linear fashion. The approximation makes the random forest more parsimonious regarding the number of parameters and thus more easily interpretable.

The fact that multiple linear regression with Lasso did perform worse on the testing set than multiple linear regression without regularization, indicates that all features add at least a little bit of information to the prediction.

Because of the connection of the used ground truth alternatives and the assumed real ground truth values (compare table 6) it may also be possible to infer CSC scores with the help of the employed algorithms.

In order to use the predicted values as CSC scores, the prediction has to be transferred back. In this step it can also be scaled to an appropriate value that is consistent with the other assessment mechanisms of the system (compare also section 9.2.9).

Part IV

INVESTIGATIONS OF THE MARKET SYSTEM

The second pillar of the contributive social capital system is the continuous assessment of CSC scores with the help of market interactions, following the concept presented in chapter 7. To investigate this concept, several experiments were conducted. Chapter 13 discusses related market mechanisms and sets them in context to the CSC market. In chapter 14, an experiment is described that was conducted to test how the CSC market system can be used to build CSC scores from currency transactions. Chapter 15 expands these investigations with the help of a market simulation with 1,000 agents over a longer time frame.

13 USING MARKET SYSTEMS FOR CSC ASSESSMENT

In our world, accumulating wealth is often seen as an indicator for economic success, as money is usually earned as wages, salaries or in interactions related to economic endeavors. As described in chapter 7, the social capital market aims to correlate success in the market with activities that increase the overall social capital of society, e.g., by providing help and information for others. The market is consequently used as a complex CSC management system that rewards helpful contributions by employing currency as feedback token. The theoretical concept of the market is described in chapter 7. In this chapter, the market system is set in relation to other feedback systems (section 13.1). Different (online) markets and related concepts are reviewed in section 13.2 and set in relation to CSC markets. The discussion is relatively brief, as there is, to the best of our knowledge, no previous work on the assessment of CSC or related properties with the help of market systems. Finally, section 13.3, classifies the CSC market in relation to other markets.

In the following two chapters 14 and 15, the social capital market is investigated directly with two experiments.

13.1 FEEDBACK SYSTEMS AND THE CSC MARKET

In section 3, social networking platforms, microblogging, scientometrics, and threaded discussion and Q&A portals were discussed as data sources for CSC analysis. The respective platforms often implement intrinsic feedback mechanisms. Cheng and Vassileva conclude about these feedback systems: "The main idea is to measure and reward the desirable user activities and compute a user participation measure, then cluster the users based on their participation measure into different classes [...]." (Cheng and Vassileva, 2006)

The market system we proposed as an additional feedback system in section 7, also leverages crowd-sourcing to collect information about individual participants. It is, however, not based on votes or ratings but on market interactions in which virtual currency tokens are exchanged. Nevertheless, the underlying idea is similar, as the SCC transactions can be seen as a feedback by individuals and the created CSCWs can be seen as reflection of a person's reputation regarding CSC within the system.

The openness of crowd-sourcing approaches can be an open door to exploit the system for one's own benefit. Lu at al. propose to address attacks on social welfare and other antisocial behavior with incentive mechanisms in crowd-sourcing (Lu et al., 2017). Dellarocas and Sang et al. review web-based feedback, as well as trust and reputation systems for online service provision (Dellarocas, 2003) (Sang et al., 2007).

Before we investigate how the proposed market system can be used as a CSC feedback mechanism in the next chapters, we briefly review related market systems in the following sections.

13.2 MARKETS AND MARKET MECHANISMS

Markets are places that allow buyers and sellers to interact based on rules of demand and supply. There is a variety of markets that range from physical consumer and business markets to online and financial markets. Sometimes market systems are even used to create forecasts or to incentivize wanted behavior. While some markets allow barter, most markets use money as exchange token. Markets are an integral component of everyday live and most goods and services are bought in markets. (Aspers, 2011)

A similarity between the CSC market and other markets is that the latter can to some extent be described as feedback system as well. With every buying decision customers indirectly cast a vote for which products and services are important to them and which suppliers they want to support. An example is groceries shopping in the supermarket. Customers can decide whether they want organic or regular products and the wholesaler and producers will often adjust their portfolio, depending on what the customers buy.

In the following, the most prominent markets are presented and set in relation to the second pillar of the CSC system.

13.2.1 *Physical Markets*

The most popular examples of markets are physical consumer and business markets. Consumer markets offer products directly to the final customer and include retailers, marketplaces, big box stores like supermarkets and also auctions, flea markets, and fairs. Business markets describe the sale from business to business. Typical examples are wholesale markets, trade fairs, and auctions, but also labor markets where people can sell their labor in exchange for wages. The price finding often takes place indirectly via price comparisons and adjustments but may also take place via negotiations. (Aspers, 2011)

The CSC market is not designed as a physical market. People may still interact and exchange services or goods offline, the infrastructure and payment options, however, are exclusively online.

13.2.2 *Virtual and Online Markets*

Some online markets like eBay and Amazon are online realizations of physical markets that obey the same principles. In this section we review three examples of other types of online markets that are related to the CSC system: cryptocurrency exchanges, information and privacy markets, and virtual in-game markets.

Cryptocurrency exchanges may appear to be virtual because all interactions take place online. The currency however, is closely linked to other major currencies like the USD or EUR and can – in the case of Bitcoin – even be traded as future at the CME and CBOE stock exchanges. Cryptocurrency can be used to buy or sell real products and services, even though it is still mainly used as investment vehicle. (Tapscott and Tapscott, 2016) The underlying blockchain technology provides security for the transactions and may be leveraged for the social capital market system, as described in section 19.3. The direct exchange of SCC for other currency, however, is not planned to prevent an intrinsic advantage for wealthy people.

Another interesting concept was presented by Groh and Birnkammerer: Information and privacy markets. In (Groh and Birnkammerer, 2011), they introduce an XACML-based approach that allows to control the information flows on social networking platforms. As the amount of shared information is often not actively controlled by the users of OSNEM platforms, the platform providers often 'own' many aspects of the associated information flow (Groh and Birnkammerer, 2011). The value of this information becomes apparent when looking at the market capitalization of companies like Facebook. Information and privacy markets are an attempt to put the platform user back in the driver's seat and allow them to actively control (1) which rights of an information item they produced are offered to whom, and (2) which information items they wish to consume. By defining these needs it becomes possible to trade information items in a market-like setting, which may lead to an increase of privacy awareness and even foster information quality. This approach is partly related to CSC markets in which information items are also one of the main goods that are traded. In CSC markets, users are also put in the driver's seat: they can buy and sell information in exchange for SCC, instead of being restricted to a pure consumer position.

Another online market that is truly virtual is realized in the context of computer games. Some games provide a simulated version of a reality which players can explore. By completing tasks, they may uncover virtual items that improve the gaming experience. The popular online game World of Warcraft allows players to trade such items with each other. There is a virtual currency to facilitate the trading. This is a purely virtual market that, in some instances, may be related to real markets. Rare items are regarded as valuable and sometimes real currency is exchanged before a trade of the virtual good happens. This is an effect that may also take place in the CSC system. Once the social capital currency is established and people perceive the increase in CSCW as real value-add, it is possible that it will be traded unofficially in exchange for USD or EUR. This would reflect the usefulness of SCC. Policing and sanctioning (section 9.2.7) are required to prevent this from happening on a large scale to maintain the decoupling of SCC and real-world currency.

13.2.3 *Financial Markets*

All types of liquid assets can be traded in financial markets. There are stock, bond, currency, and future markets that can be used to trade company shares, government issued bonds, different currencies, futures, options, and other related financial products. (Davidson, 2003)

The actors in financial markets often act as buyers and sellers, which is similar to CSC markets. The traded products and the underlying idea are different.

13.2.4 *Prediction Markets*

Prediction markets are virtual markets that try to create forecasts of events by allowing actors to trade contracts that are related to the respective outcomes. The market price of each possible outcome can be seen as indicator for the likelihood of that outcome. In this instance prediction markets are related to betting exchanges or stock exchanges.

Wolfers and Zitzewitz describe three main contract types in prediction markets: winner-take-all, index, and spread contracts. (Wolfers and Zitzewitz, 2004)

Winner-take-all is a contract that pays, e.g., 1 USD if a event y occurs, e.g., that a certain candidate wins the popular vote in the US elections. If the candidate looses, the contract pays nothing. Before the event occurs the contract price reflects the estimated probability of y happening.

The *index* contract reflects the mean value of an outcome as expected by the market participants. In the case of an election the contract could, e.g., pay 1 USD for every percent point the candidate receives in the popular vote. The contract price, a value between 0 and 100, reflects the expected outcome in percent.

A *spread* contract can be used to recreate the median of the expected outcome. It may, e.g., cost 1 USD and pay 2 USD if the candidate received more than y% of the popular vote and 0 USD otherwise.

Users will consequently buy and sell according to their own expectations. Thus, the market price reflects a community assessment. Wolfers and Zitzewitz showed that "[...] market-generated forecasts are typically fairly accurate, and that they outperform most moderately sophisticated benchmarks" (Wolfers and Zitzewitz, 2004).

CSC markets do not work like prediction markets as they do not offer to buy or sell contracts about other person's CSCWs. They aim at recreating a person's CSCW with the help of direct SCC transactions. Because of the differences, the success of the prediction markets is no guarantee that the CSC market assessments will also be accurate. Nevertheless, it underlines the usefulness of community assessments and motivates the CSC market research.

13.2.5 *Market Systems for Governance*

An example for a governance system that is based on market principles was the decentralized autonomous organization (DAO) called "The DAO" (Buterin et al., 2014). The DAO was a distributed system that was operational in 2016. It was implemented as smart contracts on the Ethereum blockchain and was consequently fully transparent. Users could buy shares of the system and would obtain the currency "DAO tokens" in exchange. At the height of the system over 150 million USD were invested in the DAO.

The DAO was conceptualized as a hub that disperses money into commercial and non-profit enterprises. Users could vote for proposals with a weight that was related to their DAO tokens. This direct involvement made conventional management structures obsolete. Profits of the enterprises were then led back to the shareholders.

The selection approach is related to the community assessment of prediction markets and can be described as flat shareholder democracy.

Due to a flawed implementation, the DAO was hacked in June 2016. Even though all funds were restored, the system never recovered.

The CSC market does not aim to be a "hub that disperses funds". The decentralized structure, the transparency, and the empowerment of individuals, however, are shared commonalities.

13.2.6 *Market Systems to Incentivize Behavior*

Market mechanics can also be used to incentivize behavior, e.g., to regulate energy and power consumption (Spees and Lave, 2007) or to promote sustainable tourism (Mycoo, 2006). Another popular example is the European Union Emissions Trading Scheme (EU ETS) that was launched in 2005 to reduce CO_2 emissions (Ellerman and Buchner, 2007).

The underlying idea is to make unwanted behavior more expensive. In the case of the EU ETS this was achieved by introducing a new commodity, the right to emit CO_2. The amount of CO_2 that can be emitted is restricted for all participants and the emission rights are auctioned off or allocated for free. Large emitters of CO_2 must monitor and report their emissions and hand in a sufficient amount of emission allowances. This creates an economical incentive to reduce the emissions in order to prevent the need to buy more CO_2 allowances.

While the European Union reports that the EU ETS reduced CO_2 emissions by an average of 8 percent (EuropeanCommision, 2010), other reports indicate that the reduction is only marginal (Anger and Köhler, 2010).

The second pillar of the CSC system also implements a limited resource – in this case SCC – which can be traded among the participants. Its limitation provides value to the SCC and motivates people to spend it primarily for goods that they really want, e.g., for interesting information items. This prioritization is similar to the EU ETS market and to prediction markets and may allow for a CSCW assessment by the community.

13.3 CHARACTERIZING THE SOCIAL CAPITAL MARKET

In the last section different markets were reviewed. Following (Aspers, 2011), there are several ways to characterize markets. To create a better understanding of the second pillar of the CSC system, this section discusses these characteristics with regard to the CSC market.

A market can be defined as "[...] a social structure for the exchange of rights in which offers are evaluated and priced, and compete with one another." (Aspers, 2011)

The *social structure* describes the setup of the market in which buyers and sellers can interact and trade with each other. In the case of the CSC market, this platform is online. Buyers and sellers may still interact offline, but the transactions and the build up of CSCW take place within the online system.

The *evaluation, pricing, and competing* is not restricted by the CSC system. All participants can freely decide what to sell and buy and at what price.

Aspers introduces two additional typologies for the characterization of markets. The first ordering typology relates to the role of the actors in the market. There are *fixed-role* markets, in which sellers and buyers are always sellers and buyers, and *switch-role* markets in which sellers and buyers may switch places (Aspers, 2011). An example for the former is a wholesale market, for the latter the stock exchange.

The second typology is whether the market is a *status* or *standards* market. This mainly relates to the product. The quality of a product that is standardized can be compared independently of the transaction partner. An example for standardized products are again shares bought at the stock market. As a share of one company is clearly defined, the identity of the individual seller is usually not important. The opposite of standardized products are status products. An example is fashion, which is highly dependent on the respective brand or designer. (Aspers, 2011)

With these two topologies a 2x2 matrix can be created to classify markets. In the CSC market the roles of the actors are not fixed. Each participant can act as seller and buyer. The provided information items or services are usually related to the seller, i.e. a news article by the New York Times is likely perceived as more valuable as an article by a private, previously unknown individual. The CSC market can consequently be classified as a *switch-role status market*.

Another example for a switch-role status market is the (oriental) bazaar (Aspers, 2011). It is made up of actors who often do both, buy and sell. The market is not focused on a standardized item but on many different goods. While some products are certainly comparable, the reputation of the sellers plays an important role.

We expect that the goods and services traded in the CSC market vary widely from those sold at an oriental bazaar and are more focused on information items and other "virtual" goods due to the non-exchangeability of SCC to real-world currencies. However, the basic principles regarding initiating a trade and price finding are similar. The exceptions are that the currency is different and that the CSC market directly interprets transactions as micro endorsements to create CSC scores. These two characteristics are more related to online feedback mechanisms, as discussed in section 13.1.

14

EXPERIMENT TO INVESTIGATE CSC ASSESSMENT WITH MARKET SYSTEMS

14.1 SYNOPSIS

The second pillar of the CSC system is the continuous assessment of CSCWs by other participants. Chapters 5 and 7 argued that a market system with an intrinsic CSCW assessment mechanism might be a suitable way to realize this goal. This idea is supported by the success of prediction markets which indicates that markets can be used for more than mere economic transactions. (Wolfers and Zitzewitz, 2004) In chapter 13, other markets were discussed and set in relation to the CSC market. To the best of our knowledge, there is no work that directly employs a market system for the assessment of CSC or related characteristics in the way the CSC market does.

To investigate to what extent markets can be employed in the context of the CSC system, an experiment with 242 participants was conducted. This investigation was part of the experiment presented in chapter 11 and allowed users to trade a virtual currency for help or information. Following the CSC build-up mechanism, several CSCWs were computed for each participant based on their transaction history. The calculated contributive social capital weights correlate with the respective ground truth peer assessments, which may be seen as an indication for the suitability of markets for the CSC assessment.

14.2 MOTIVATION AND RESEARCH QUESTIONS

Using a market system for the continuous assessment of CSC brings several advantages that are in line with the goals of the system:

- The assessment happens automatically based on the transactions — no additional work on the side of the participants is required.
- The existence of currency may allow new business models.
- People may be motivated to provide services to the community as they can receive rewards in the form of SCC transactions.

The experiment was guided by the following research questions:

1. What is the perception of the market by the participants?
2. Do the CSCWs calculated with the help of the market system correlate with the ground truth assessment?
3. How can the CSCWs be divided along different topics?
4. Does a basic income of virtual currency lead to an inflation within the system?

14.3 DESCRIPTION OF THE EXPERIMENT

The experiment was conducted over a period of nine weeks, during the summer semester 2017, with 242 volunteers from the social computing lecture at Technical University Munich. The participants also took part in the experiment that investigated the extraction of CSC from online social networking platforms, as described in chapter 11. This allowed us to use a similar infrastructure and the same methods for motivating a fair and balanced participation. It also allowed the participants to engage on the social networking platform and form an opinion about other participants based on these interactions. Three topics that were heavily discussed on the platform were "Populism in politics", "Living in Munich", and "Healthy food and sustainability" (see appendix B for a list of exemplary discussion topics).

With the help of the python-based Django[1] framework, a web portal was created that allowed the users to participate in the market. A screenshot of the interface is provided in appendix C. In line with the discussions in chapter 7, the market system created for the experimental investigation, features the following functionalities:

- *Uniquely identifiable profiles:* Registration with a profile picture and name or a selected pseudonym. This profile is uniquely linked to the ground truth assessment.
- *Currency distribution algorithm:* 100 SCC were handed out at the time of registration and an additional 100 SCC as basic income every week.
- *Transactions between users:* For every transaction the participants could specify the recipient, an amount between 1 and 100 SCC, one of six categories, and a personal message. There were three categories related to the knowledge topics ("Populism in politics", "Living in Munich", "Healthy food and sustainability") and the categories "Trustworthiness", "Social responsibility", and "Other".
- *CSCW build-up mechanism:* Every transaction was used for the assessment of the recipient's $CSCW^{i}_{\mathtt{recipient}}$ in topic i following equation 27. The factor α was chosen to be 1.005, which allowed a visible increase for smaller transactions while preventing huge increases for larger SCC transfers.

The students were asked to participate in the market in whatever way they liked, e.g., to reward and acknowledge helpful contributions on the social networking platform, or to pay for information or help they received from others. The market could be accessed at any time during the nine-week duration of the experiment.

14.4 DESCRIPTION OF THE CREATED DATA SET

The demographics of the participants are described in the social network analysis part of this experiment, in section 11.4.2.

894 transactions with an average value of 42.4 SCC were completed over the course of the experiment. The count of transactions and the highest CSCW values by category are provided in table 15.

1 https://www.djangoproject.com/ (retrieved: 2018-05-22)

Category	Count of Transactions	Highest CSCW
Populism in politics	58	2.7
Living in Munich	149	3.0
Healthy food and sustainability	73	1.8
Trustworthiness	148	4.9
Social responsibility	121	7.4
Other	345	9.0
Total	894	-

Table 15: Count of transactions and highest CSCW values by category

The most transactions were labeled "Other", followed by the topic "Living in Munich" and general "Trustworthiness". The least amount of transactions were labeled "Populism in politics", even though there was no significant difference regarding the number of posts about politics and living in Munich on the social networking platform. This may indicate that the topics related to life in Munich were perceived as more helpful than the discussions of politics.

The column "Highest CSCW" lists the highest weight that an individual achieved based on received transactions.

The highest social capital score was achieved by the student with the user name *Bob*[2] in the category "Other". Thanks to the support of 24 other students who transferred SCC to *Bob*, he built a CSCW of 9.0. Due to the smaller amount of overall transactions, the maximum weights in the other categories are smaller. The highest value in the category "Social responsibility" was achieved by *Alice*, who also scored highest in the category "Healthy food and sustainability". The trustworthiness score of 4.9 was also achieved by *Bob*. *Carol* holds the highest CSCW in "Populism in politics" and *Dave* received the highest score in the topic "Living in Munich". The personal messages in the respective transactions indicate that the transactions were usually inspired by helpful contributions on the social networking platform. *Dave*, e.g., received SCC in the category "Living in Munich" from nine different participants who all praise his contributions and his recommendations (e.g., "Best coffee in Munich", "Secret attractions of Munich", etc.). All of these four students were rated higher than average in the respective ground truth assessments.

As the market investigations were part of the social capital experiment (see chapter 11), the same CSC ground truth, which was collected via 539 peer assessments, could be used. Of the 242 participants a ground truth assessment was provided for 165 students (average of 3.3 assessments per person). The data described in this section, as well as the data used for the following analyses was created by these 165 participants.

Additionally, a questionnaire was provided that the participants could answer voluntarily to describe their experiences and provide feedback.

2 The user names mentioned in this section were pseudonymised to further protect the identities of the participants.

14.5 CSC ANALYSIS

Based on the collected data and additional questionnaires we can draw several conclusions.

14.5.1 *User Perception of the Market System*

There are several indicators that the market was perceived as useful and sustainable. The number of transactions displays uninterrupted activity over the course of the experiment (figure 14). The decrease in week 4 is due to the fact that there was no lecture due to a public holiday, which put the experiment on pause for some days. The decline in week 9 can be explained by the end of the experiment in the middle of a week. The total number of transactions (894) is about a quarter of the 3,651 likes that were exchanged on the social networking platform. This is a reasonable ratio, considering that making a transaction is more time-consuming than "liking" a contribution.

115 participants answered a questionnaire about their experiences. 64% agreed with the statement that "Social capital currency is valuable because of its limited availability" and 61% stated that "When I received social capital currency it felt better than receiving a "like" because of the possibility to re-use it." The decision to decouple SCC from real-life currency was questioned by 40% of participants who supported the statement "If the system is rolled out on a large scale, people should have the possibility to buy/sell social capital currency with real currency." 73% of the participants experienced some kind of inflation over the course of the experiment. More than half of the participants (54%) said they would use SCC to pay for news online, if it were a possibility and 57% said they were satisfied with the amount of SCC distributed each week.

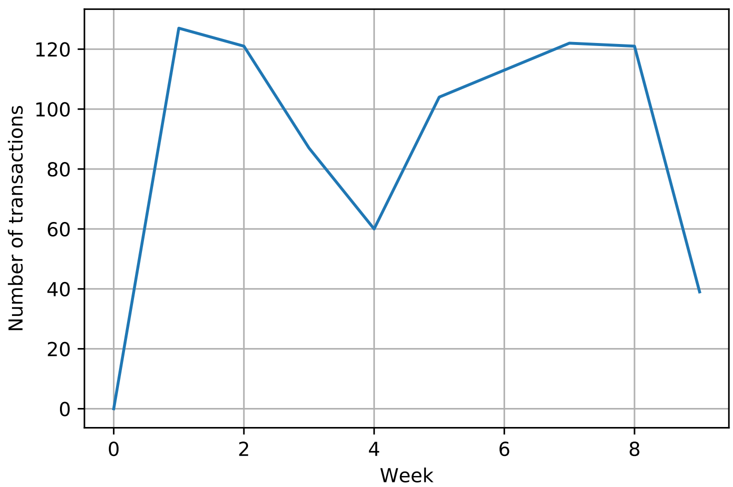

Figure 14: Number of transactions per week in the market experiment

Table 16: Spearman correlation between market assessment and ground truth value by category. The highest value by column and row is highlighted in green.

	Ground Truth				
Market Assessment	$GT_{Pop.}$	$GT_{Mun.}$	$GT_{Food/Sust.}$	GT_{Trust}	$GT_{Respo.}$
$MA_{Populism}$	0.40	-0.05	0.06	0.03	0.00
$MA_{MunichLiving}$	0.09	0.35	0.17	0.14	0.28
$MA_{Food/Sustain.}$	0.06	0.11	0.41	0.15	0.14
MA_{Trust}	0.05	-0.04	0.08	0.21	0.07
$MA_{Responsibility}$	0.09	0.11	0.14	0.10	0.26

14.5.2 *Correlation Between the Market-Determined CSCWs and the Ground Truth Assessment*

To test the hypothesis that social capital markets can be used for CSC assessments, we calculated the correlation coefficient between the topic-dependent CSC weights assigned by the market and the corresponding ground truth assessments for everyone for whom a market prediction was available. The Spearman correlation coefficients are visualized in table 16. The largest values for each column and row are highlighted in green. With one exception, the values on the diagonal are the largest. That means the highest correlation can be found between the ground truth and the corresponding market assessments. The only off-diagonal correlation value that is higher than 0.20 is between the ground truth assessment of "Social responsibility" and the market assessment for "Living in Munich". A reason for this could be that people who were perceived as socially responsible, also provided tips about Munich in order to help others. The largest correlation values (between 0.35 and 0.41) are in the three knowledge categories "Populism in politics", "Living in Munich", and "Healthy food and sustainability". A possible explanation is that knowledge and expertise are easier to assess than the less clearly defined concepts of "Trustworthiness" (0.21) and "Social responsibility" (0.26). All correlations on the diagonal are statistically significant on the 0.001 level (table 17) and thereby support the hypothesis that CSC can be modeled with the help of market transactions.

14.5.3 *Development of Buying Power over Time*

We also investigated the development of buying power over time, as the majority of participants experienced a devaluation of the distributed currency (see section 14.5.1). The average size of all weekly transactions increased for all topics, which is illustrated in figure 15. The average increase was 14% per week. This supports the impression of the participants that more SCC was required as time progressed. Due to the basic income of 100 SCC per week, the overall available currency in the system increased on average by 34% per week. This might indicate that the perceived inflation, which is reflected by the average transaction size, is smaller than the actual increase of

Table 17: P-values for the Spearman correlation between market assessment and ground truth value by category.

	Ground Truth				
Market Assessment	$GT_{Pop.}$	$GT_{Mun.}$	$GT_{Food/Sust.}$	GT_{Trust}	$GT_{Respo.}$
$MA_{Populism}$	<0.001	0.525	0.418	0.735	0.995
$MA_{MunichLiving}$	0.247	<0.001	0.034	0.073	<0.001
$MA_{Food/Sustain.}$	0.473	0.169	<0.001	0.056	0.077
MA_{Trust}	0.55	0.653	0.303	0.008	0.366
$MA_{Responsibility}$	0.231	0.165	0.079	0.186	<0.001

available currency. The variance in the topic "Populism in politics" can be explained by the low number of transactions in week 4 and week 9.

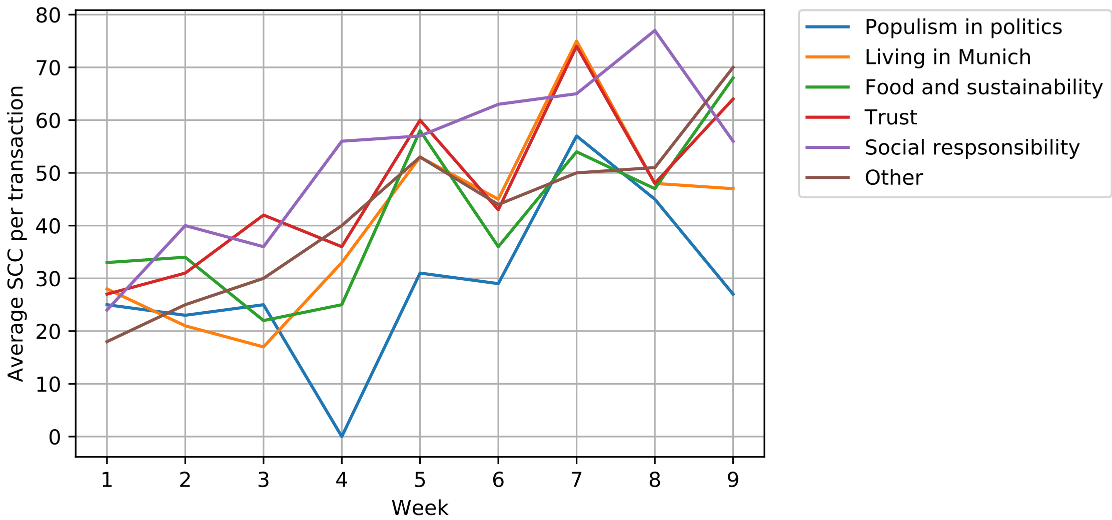

Figure 15: Development of SCC transaction sizes per week over the course of the market experiment

14.6 POTENTIAL SHORTCOMINGS OF THE EXPERIMENT

The market investigations were conducted together with the social capital experiment. The shortcomings of the experiment are described in section 11.6. An additional shortcoming may have been the novelty of the market system, which may have led to initial unnatural behavior, e.g., to "test the system".

14.7 DISCUSSION OF RESULTS

With these limitations in mind we can draw several conclusions from the analyses based on the five initial research questions stated in section 14.2.

The first research question was about the feasibility to implement a market that is perceived as useful by the users. The python-based market infrastructure ran without problems on a university server. No one reported any performance problems, even though it was used simultaneously by many students. The continuous usage of the system and the positive feedback (as described in section 14.5.1) further support the hypothesis that such a social capital market can be designed and successfully maintained. For an implementation on a larger scale, e.g., for future experiments with several thousand participants, we recommend two improvements. Firstly, a more effective searching algorithm. In this experiment, users could identify each other with their profile picture and search for the pseudonym. This is sufficient for about 200 users, but for larger numbers, a more detailed search may improve the user experience. Secondly, we experienced about 15 students forgetting their passwords. These passwords had to be reset manually in the database (where a hashed version was saved), which was time-consuming. If a similar percentage of participants forget their password in a larger experiment, it is essential to implement an automatic reset functionality.

Research questions 2 and 3 asked for the possibility to infer (topic-specific) CSCWs from the market interactions. Thanks to the implemented mechanism, which is described by equation 27, such an CSCW assessment happens automatically. The experiment investigated this assessment on an individual level and on the whole data set. The four students who achieved the highest scores in the six different categories, were also rated better in the respective ground truth assessments. By correlating the lists of users ranked by the ground truth peer assessments and the market predictions, we investigated the hypothesis on the whole data set. As discussed in section 14.5.2, the correlation between the related lists is in general much higher than between unrelated lists. The correlation indicates that it is possible to identify CSC with the help of markets to a certain extent and that the CSC can be measured on the level of individual topics. The contextualization by topics was achieved by providing different categories for the exchange, which allowed for a build up of different CSCWs.

These findings are important for the fourth and fifth overarching research question (section 1.2) as they underline the feasibility of a market-based assessment. This is also in line with previous work by Wolfers and Zitzewitz who demonstrated that market-based community assessments could be used to create fairly accurate forecasts with the help of prediction markets (Wolfers and Zitzewitz, 2004).

The fourth research question regards the distribution mechanism of the social capital currency. This mechanism is required to allow everyone to participate in the system. The amount of the basic income was perceived as sufficient by a majority of the users who answered the questionnaire (see section 14.5.1). However, 73% experienced some kind of inflation, due to the increasing amount of available SCC. This subjective impression was supported by the analysis of the average transaction size (figure 15). If this is unwanted in a future implementation of a market system, it can likely be reduced by the introduction of either sales or capital taxes, as discussed in section 9.2.5.

The current design of the system allows fraud in the form of *pumping* CSCWs by transferring SCC back and forth between two or more users without valid reason to do so. We did not find any traces of systematic pumping in the database. For a larger scale experiment the possibility of pumping should, nevertheless, be reduced by either transaction taxes or suitable detection mechanisms (compare discussions in sections 9.2.5 and 9.2.7).

To investigate the behavior of the market on a larger scale, we ran several simulations. These experiments are described in the following chapter 15.

15

SIMULATION OF A MARKET SYSTEM

The simulation experiments were created in the context of the bachelor's theses of Maximilian Schmidt (Schmidt, 2017) and Niclas Hirtle (Hirtle, 2018) that were both supervised by Sebastian Schams and Georg Groh in the Social Computing Research Group at Technical University Munich.

15.1 SYNOPSIS

Chapter 14 presented an experiment that investigated how a market system can be used for the assessment of individual and topic-specific CSC scores. A major shortcoming of this experiment was its scale; the CSCW could only be assessed of 165 participants who participated in the market during a time period of nine weeks. To address this deficit, we conducted market simulations. With the help of these simulations it was possible to investigate the quality of the assessment, inflation within the system, and other parameters on a larger scale. The findings of this chapter further support the hypothesis that the ability of markets to assess individual properties like CSC are worth investigating.

15.2 MOTIVATION AND RESEARCH QUESTIONS

Simulations can be understood as "models which show the working of very complex relationships, those which are too complex to be reduced to simple conclusions by means of mathematical or statistical analysis, or ordinary reasoning" (R. Verma, 2010). The interactions of a market, especially over a long time period, are nearly impossible to describe with statistical analysis or ordinary reasoning alone. Consequently we leveraged the possibilities of *Agent-based Modeling and Simulation* (ABMS) to further investigate interactions in a social capital market system. ABMS is generated based on the actions and interactions of individual agents and is widely used for the analysis of complex systems in different fields of research (Klügl and Bazzan, 2012). We use it to investigate the following research questions:

1. Does the correlation between the true CSCW and the CSCW that is assumed to build-up during the interactions, improve when the experiment is run over a long time?

2. Is there inflation present and how does the average transaction price behave over time?

3. After what time does the simulation achieve a satisfying depiction of the real CSCWs?

4. How is the SCC distributed among the agents at the beginning and end of the experiment? How does this compare to income distributions in real countries?

5. What scaling options should be used to represent CSCWs? In the previous experiment (chapter 14) no scaling was required as the transaction count was so low that the maximum scores never reached 10.

6. Which of the agents' properties had the largest effect on the prediction of CSCW?

The first two questions are similar to some of the research questions of the market experiment (chapter 14). The last four questions go beyond the investigations of the previous chapter. These questions are either simulation specific (e.g., question 6) or facilitated by the larger sample size and duration that allows insights into how the system behaves after several years.

15.3 ASSUMPTIONS AND SETUP OF THE SIMULATION

The simulation was written on Python. Each simulation was run with 1,000 agents over a simulated time period of 5 years. In real life, there is a multitude of different interactions that lead to a transfer of SCC, e.g., providing help or information, as well as rendering services or making other kinds of pro-social contributions. As the nature of these interactions is of little relevance for the simulation, we call it exchange (buying or selling) of *value*. In the simulation, the $CSCW_{real}$ grows proportionally to the accumulated *value* to reflect that the bought information item increased the CSC of the buyer. In real life, the buyer's increasing $CSCW_{real}$ can, e.g., be thought of as the person's increasing knowledge or competence. This takes the nature of the CSC market into account, which is likely focused on the exchange of virtual goods and information items, as described in section 13.3. An agent's behavior within the model is governed by several assumptions and parameters that reflect character traits which are of relevance for social capital related interactions:

- The initial value of $CSCW_{real}$ is drawn from a truncated normal distribution with mean = 4 and a minimum value of 1. $CSCW_{real}$ can change by buying *value* and represents the ground truth CSCW value that the market system aims to identify through interactions. The distribution was chosen because the CSCW ground truth assessments of the empirical experiment are distributed normally.

- The CSCW measure within the system is $CSCW_{virtual}$ and initially set to 1 for all agents. It increases when the agents receive SCC and is visible to all agents in the system.

- To approximate a real network, all 1,000 agents are connected in the form of a Watts-Strogatz small world graph (Watts and Strogatz, 1998), with dimension 1, mean degree 4, and rewiring probability 0.5. The distance to another agent in the system is relevant for the selection of an interaction partner.

- Another assumption of the simulation is that all agents are social beings and as such interact in a way that is relevant for social capital. Each interaction increases the recipient's $CSCW_{real}$ and triggers a transaction of SCC as a matter of payment or appreciation, which in turn increases the sender's $CSCW_{virtual}$.

- The trading frequency of each agent is based on observations by Miritello et al. (Miritello et al., 2013) They describe the number of partners a person interacts with on a regular basis with a truncated normal distribution with a mean of approx. 10 and the frequency of their interactions as about two per month. Under the assumption that such interactions would be followed by a transfer of SCC in the CSC system, we draw the trading frequency of each agent from a similar truncated normal distribution. The average trading frequency is consequently a little less than once a day in effect.

- The agents have preferences with whom they interact, i.e. from whom they buy *value*. The first preference is that the agents prefer to interact with someone who they already know. This is achieved by weighing all possible trading partners with the distance in the Watts-Strogatz small world graph. The second preference is the $CSCW_{virtual}$ of the other agents, which is important because sellers with a high $CSCW_{virtual}$ are preferred.

- The price for the received *value* (the amount of SCC transferred) is determined by the current market price for a similar *value* and the agent's satisfaction with the received *value*.

- After an interaction, the agents are able to assess the $CSCW_{real}$ of their counterparts and adjust their payment accordingly. They pay less than the average price if the $CSCW_{real}$ of their counterpart was lower than the previously seen $CSCW_{virtual}$ and more if it was higher. This reflects the behavior of real people who are disappointed if the advertised quality was not met (a lower $CSCW_{real}$ than $CSCW_{virtual}$ value) or thankful if their expectations were exceeded.

- There is no intentionally malicious behavior between the agents.

In every time step, each agent is randomly assigned one of the following actions, based on their individual properties:

- Place an order for a *value* at a price that is influenced by the time it takes to create the *value* and the minimum $CSCW_{real}$ required from the seller, as well as the amount of SCC owned by the agent. The price asked by the seller is influenced by the prices of the previous transactions between other agents. This reflects all kinds of interaction in real life that follow the pattern of seeking help or information on a defined topic (e.g., playing the guitar) for a required time period (e.g., practice for 1 hour) from a person with a certain level of reputation (e.g., music teacher with 5 years of experience). In the case of information items, the required time reflects the time saved buy the buyer (e.g., the information bought by a person makes a two-hour web search unnecessary).

- Offer a *value* for sale on the market, for a price that reflects the current market price and their own $CSCW_{virtual}$. The simulation implements the current market price instead of a fixed price, to reflect the price finding mechanism in other markets. This also allows to investigate how the market price develops over time.

- Do nothing, if they do not have any currency left or if they already interacted frequently enough considering their interaction preferences.

If an agent placed an offer to buy or sell a *value* for a certain price, but no other agent placed the corresponding sell or buy offer, the agent holds their offer for a defined amount of time before adjusting it to increase the chances for a match. The time an agent waits is determined by the patience parameter ρ. With each passing time step in which no interaction occurs, ρ is reduced. Once ρ reaches 0, the agent places a new order and ρ is reset to its original value. For the simulation ρ was set to two days.

The underlying market mechanism that matches the agents and builds the CSCW scores, is based on the following assumptions:

- The matching between agents is achieved via a continuous matching algorithm that is based on the Continuous Auction System described by Posada et al. (Posada et al., 2006) The algorithm assigns agents who want to buy *value* to agents who want to sell that *value* based on the price and the agents' preferences. This is similar to the price finding mechanisms we expect in switch-role status markets.

- The transactions between agents and the build-up of CSCW follows the equations 20, 21, 22, and 27. To improve visual interpretability, the build-up of CSCW is constrained to a logarithmic scale from 1 to 10 via adjustments to the factor α in equation 27. Logarithmic scales have been used successfully in other systems for the extraction of CSC related properties (compare (Rao et al., 2015)).

- Each agent starts with 100 SCC and is rewarded a monthly basic income of 100 SCC.

15.4 DESCRIPTION OF DATA SET

After the simulated period of five years, there were 1,628,084 interactions between the 1,000 agents, which corresponds to the expected average number of 0.9 interactions per agent per day. The number of transactions by an agent is distributed along a truncated normal distribution (see figure 16) with few agents having almost no transactions and a few agents having traded over 4,000 times. The interactions between the agent's have led to an increase of their real CSCW. The new $CSCW_{real}$ distribution is visualized in figure 17. The majority of agents now have a weight above 9.5. Due to the logarithmic scaling, there is still a significant difference between those weights. To what extent the virtual weights $CSCW_{virtual}$ describe these real weights after all the interactions, is described in the following section.

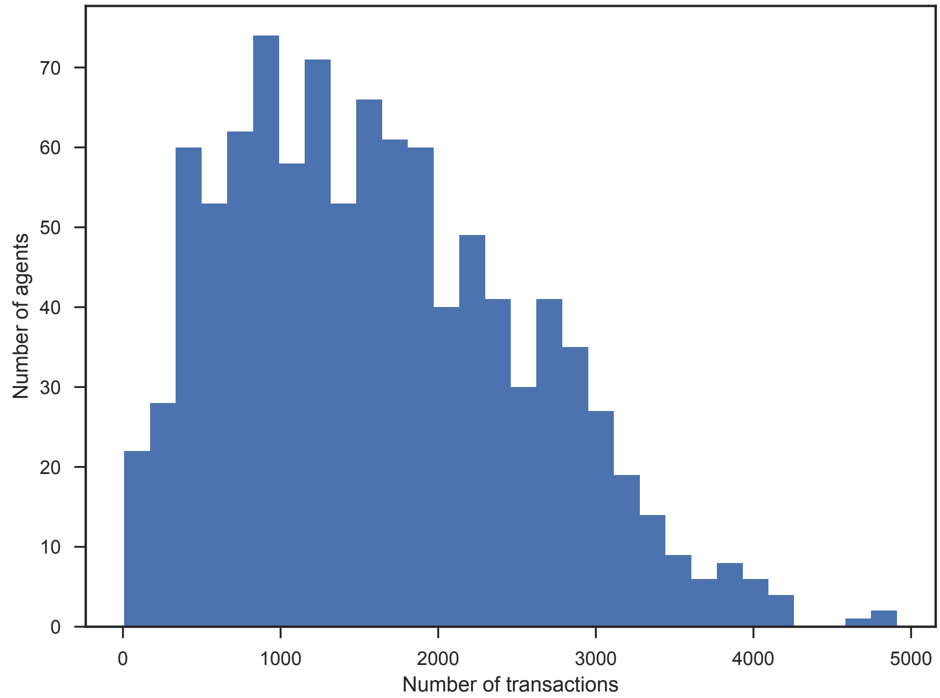

Figure 16: Number of transactions by agents at after the simulation

15.5 CSC ANALYSIS

Due to the adjustments to $CSCW_{virtual}$ with each transaction, there is now a correlation between the true and the virtual contributive social capital weights, as visualized in figure 18.

The Spearman correlation coefficient is 0.998. This correlation is facilitated by the assumption that the agents could assess each other's $CSCW_{real}$ after each interaction and adjust their payment in case they were unsatisfied or satisfied by decreasing or increasing the payment accordingly. This is a necessary requirement for the system to work. In a real world setting people are usually only able to appraise each other's CSCW to a certain degree, so it is uncertain to what extent the results carry over. However, we can make several observations based on these assumptions.

At least 30 agents lie significantly above the diagonal, which indicates that the real CSCW is under-estimated in some cases. This is especially the case for agents with a low interaction frequency, as $CSCW_{virtual}$ can only be built up through interactions. No agents lie significantly below the diagonal, i.e. few agents are over-estimated by the system. This indicates that no negative feedback mechanism is required that would allow agents to reduce the $CSCW_{virtual}$ of over-estimated agents. The accumulation of agents with high $CSCW_{real}$ and $CSCW_{virtual}$ in the upper right corner of figure 18 is due to the continuing increase of these values with each interaction and the logarithmic nature of the scales that allow only asymptotic approximation to the maximum value 10.

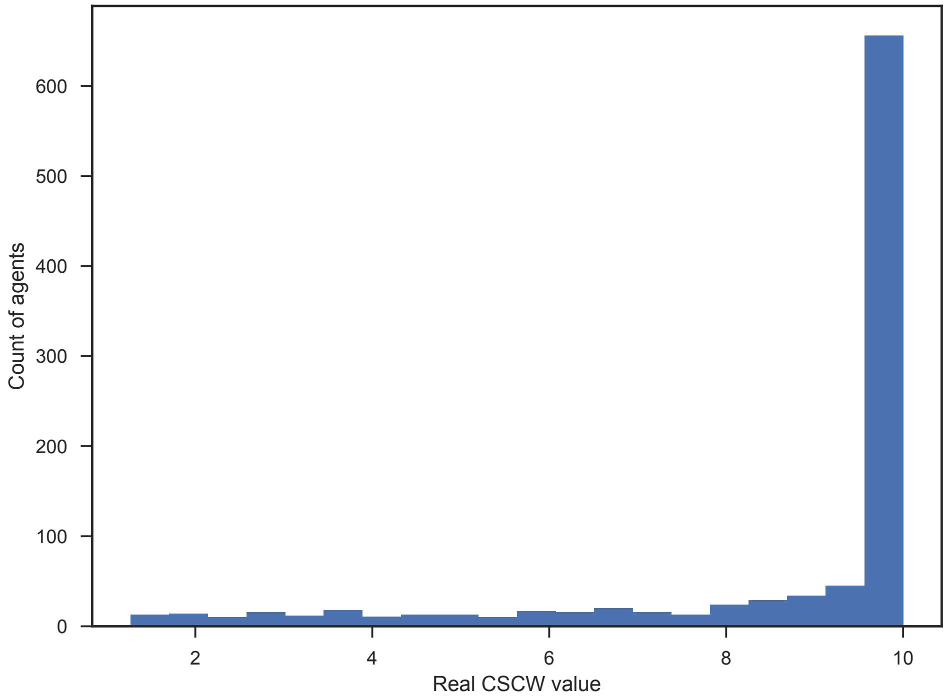

Figure 17: Distribution of the agents' real CSCW after the simulation

We are also interested in determining the time required by the system to reach a steady state, in which the majority of agents are assessed correctly. This duration is directly determined by the selected interaction frequency. As described above, this parameter originates from empirical research, so the results are meaningful even outside the simulation context. As a measure for the prediction accuracy we chose the percentage of correctly predicted agents over time by percentile (figure 19). It takes almost two years for all percentiles to reach a steady state of around 90% correctly predicted scores. The half of the agents with above average CSCW (top 50 percentile) is already predicted correctly approximately 90% of the time after one year.

With the help of the Gini coefficient (Gini, 1912), we can observe the distribution of social capital weights and SCC within the system. A value of 0 represents an equal distribution, whereas values close to 1 reflect that one or few agents pool all the resources. The results are plotted in figure 20.

The distribution of $CSCW_{virtual}$ starts from 0, as every agent has a $CSCW_{virtual} = 1$ at the beginning, and increases over the first year. After about 1 year, it asymptoticly approaches the $CSCW_{real}$ which lies around 0.4, similar to the Gini coefficient of worldwide income equality (CIA, 2016). Both lines continue to grow slowly, as people with higher $CSCW_{virtual}$ are preferred trading partners and, therefore, increase their weights faster than others.

The currency is less equally distributed. It starts from 0, quickly increases and fluctuates around 0.6, which corresponds to the Gini coefficient of the income equality in South Africa (CIA, 2016). In a real system with real participants the difference

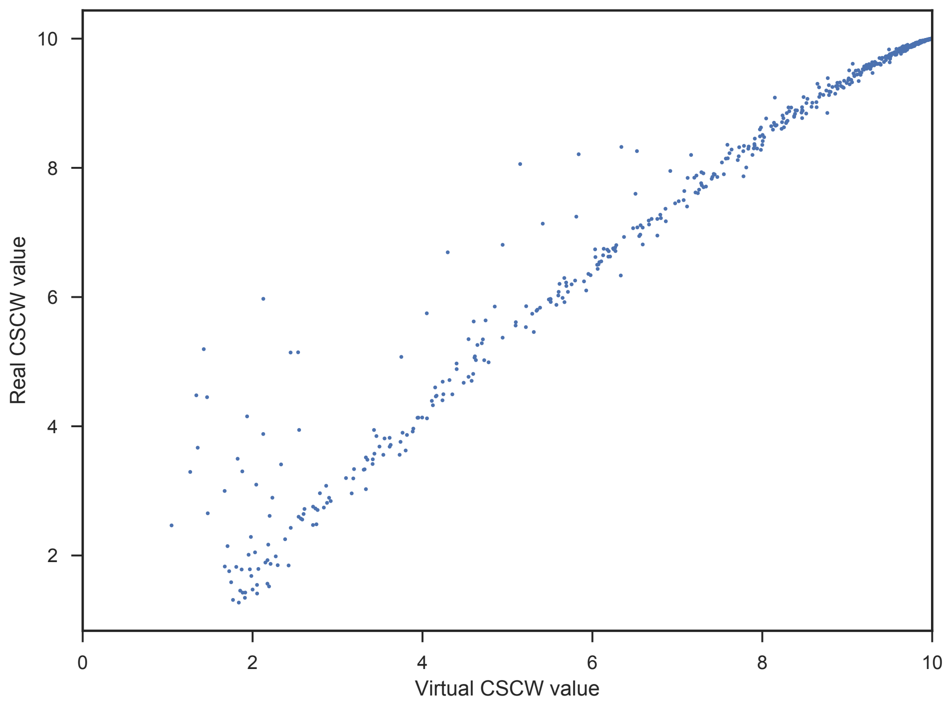

Figure 18: Scatter plot of the true and virtual CSCW after a simulated time period of five years

might not be as large, as entrepreneurial behavior by people with lower $CSCW_{virtual}$ could help to equalize the SCC distribution.

Similar to the experiment described in chapter 14, one can also see inflation within the system. The price finding mechanism makes agents bid based on their available SCC and the average market price for their desired value. The increase of available SCC due to the basic income consequently leads to higher prices with every month. This can be observed in figure 21 that displays a continuous increase of the average transaction price.

To analyze how different agent properties influence the increase of the agents' $CSCW_{virtual}$, we investigated how fast the $CSCW_{virtual}$ of different groups of agents increased. The groups were: agents with high and low interaction frequency, their betweenness centrality in the graph, and a high starting $CSCW_{real}$. Agents with a low interaction frequency had the most trouble increasing their $CSCW_{virtual}$, while the agents with the highest interaction frequency increased their $CSCW_{virtual}$ the fastest. This is due to the fact that agents need to interact and trade to receive SCC and thereby increase their weights. The well-connectedness in the graph, which was represented by a high betweenness centrality, only had marginal influence on the increase of $CSCW_{virtual}$. This is likely due to the fact that all agents had a similar betweenness centrality in our graph. The high initial $CSCW_{real}$ also helped agents significantly to increase their virtual weight. As agents with high weights are desir-

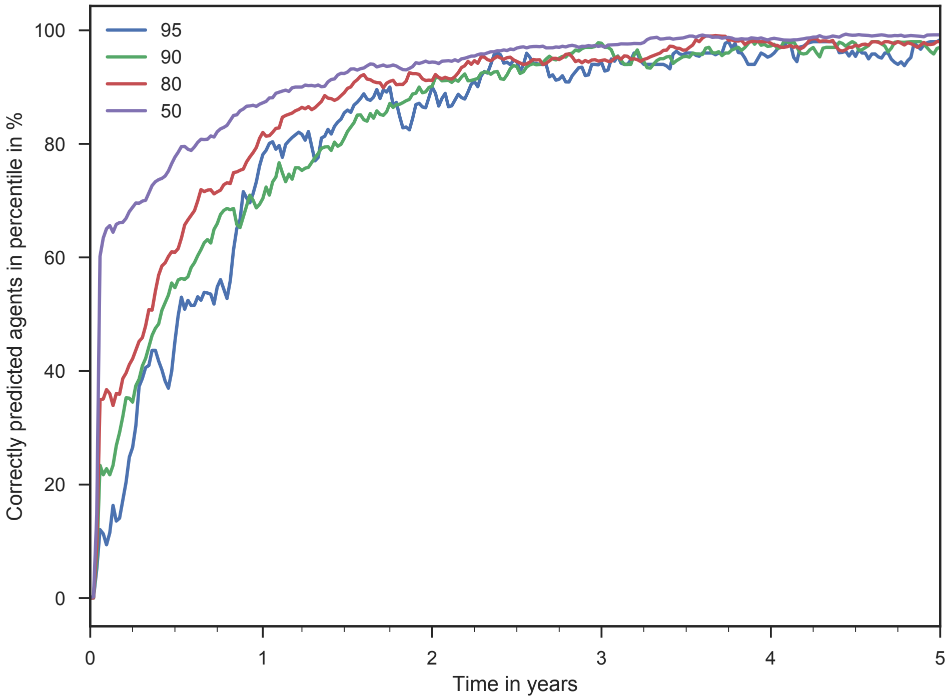

Figure 19: Percentage of correctly predicted $CSCW_{virtual}$ over time by CSCW percentile

able interaction partners, their weights increase faster than those of agents with a similar interaction preference but a lower starting $CSCW_{virtual}$.

15.6 SHORTCOMINGS OF THE MARKET SIMULATION

The ABMS of the market system is a simulation based on assumptions. The assumptions were build on experiences drawn from the market experiment (chapter 14) and other previous research. Nevertheless, there are several shortcomings:

- In the real world, agents might not be able to correctly assess each other's $CSCW_{real}$ after an interaction.
- The agents might not adjust their payments to reflect whether they are satisfied or unsatisfied.
- In the simulation, there is no entrepreneurial behavior as the agents are not intelligent and do not adjust their behavior over time.
- There are neither new users, nor drop-outs, as would be expected in a real-world setting.
- There is no distinction between topics, which led to a build up of CSCW in only one category and consequently to more experts than we would have expected if the transactions were distributed over several topics.

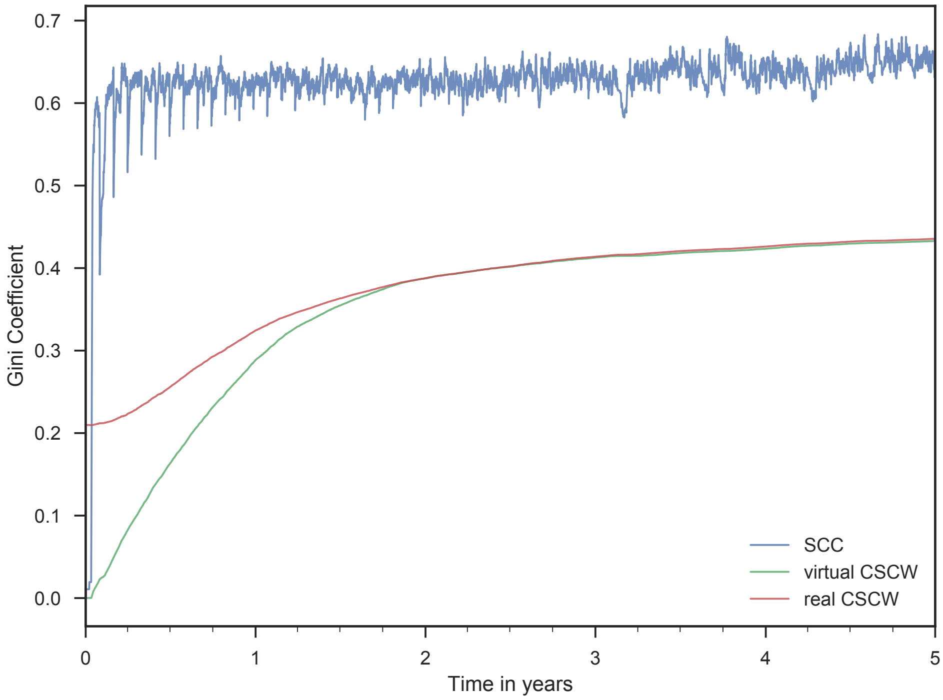

Figure 20: Distribution of wealth as represented by the Gini coefficient for SCC and CSCW over time

15.7 DISCUSSION OF RESULTS

In general, the simulation supports the findings of the social capital market experiment (chapter 14).

The goal of the simulation was to additionally investigate how the system behaves on a large scale. The first research question asked for the correlation between $CSCW_{real}$ and $CSCW_{virtual}$ after a longer time period. As visualized in figure 18, the virtual and real weights of most agents lie on a straight line. The correlation coefficient is 0.998, which underlines that the virtual CSCWs closely represent the real CSCWs.

Similar to the market experiment there was also inflation due to the basic income (research question 2). This becomes apparent when investigating the average price of a transaction over the time of the experiment. The average price varies on a daily basis, depending on the willingness and availability of agents who want to pay the current price. Additionally, the price increases by a factor of about five over the course of the simulated five year period. If this increase is unwanted, it can likely be reduced by introducing transaction or capital taxes.

It took the simulated market over a year to predict the $CSCW_{real}$ correctly in most cases (research question 3). That may appear to be a significantly longer time period than the one of the market experiment described in chapter 14. The experiment ran over the course of 9 weeks and was still able to produce a positive correlation

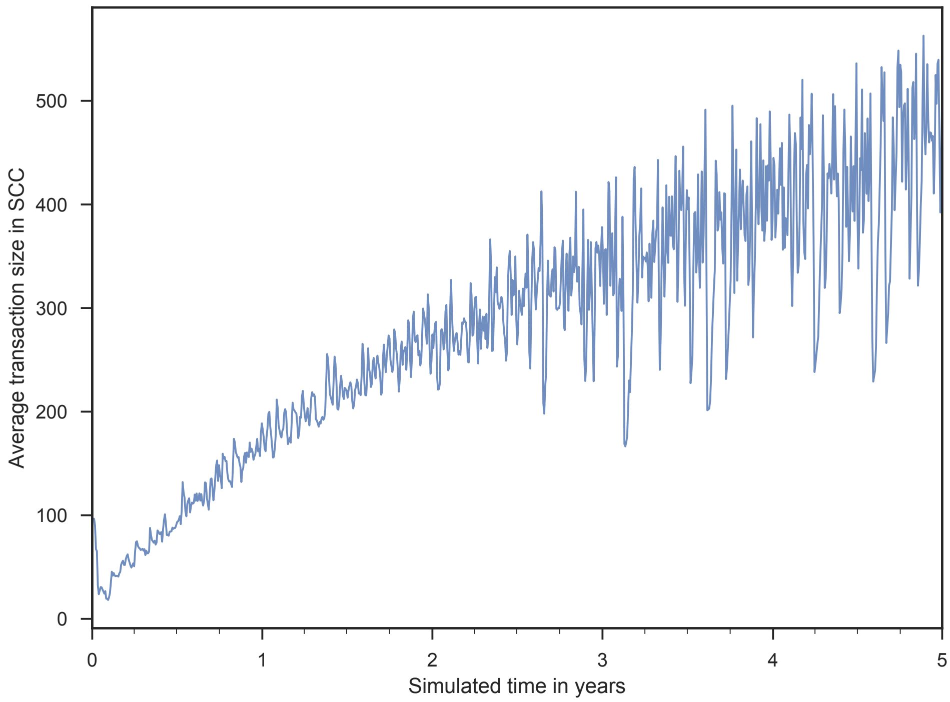

Figure 21: Price increase over time because of the basic income

between real and virtual weights. However, the number of correctly predicted CSCWs was significantly lower and the experiment was accelerated, as the basic income was handed out on a weekly and not monthly basis. The correctly predicted values of the agents with high CSCWs in the simulation is already high after much less than one year.

The SCC and CSCWs are not equally distributed among the agents (research question 4). This is to be expected due to the agent's preferences to interact with partners with a high weight. People with high CSCW consequently accumulate more SCC, as is exemplified by a Gini coefficient of around 0.6. This may be explained by the "the rich get richer" phenomenon which is described by Price's model, a generalization of the Simon model (Simon, 1955). At the hands of citation networks, Price's model describes that the probability that an existing node receives new edges is related to the number of existing edges (Fernandez-Cano et al., 2004). In citation networks the node is an existing paper and an edge is a citation. In the case of CSC markets the node is a market participant and the edges are currency transfers from new trading partners. This *cumulative advantage* of participants with a high CSCW consequently leads to a further increase of the person's weight and a CSCW that is distributed according to a power law. On the one hand this is wanted, as the system aims to reward participants who contribute to increase the overall social capital of the network. On the other hand it is important not to discourage other people from participating due to a perceived high inequality.

The CSCWs were scaled logarithmically. By running an alternative simulation with a linear scale, it becomes apparent why the logarithmic scale is preferable (research question 5). While the correlation between $CSCW_{real}$ and $CSCW_{virtual}$ still works in the linear simulation, the weights of individual agents reach values of over 100,000. Such large numbers may be impractical for every day use and contradict the goal to allow quick judgements. The logarithmic scale we employed restricted the weights to a value between 1 and 10. The investigation of the agent's weights after the simulated period of five years reveals that many of them have a CSCW above 9.5 (see figure 17). Such a value corresponds to an expert who has invested a significant amount of time into the topic. As the agents in the simulation only traded in one topic, many of them actually became experts because of the *value* they bought from others. In a real world setting, where different topics are available, it is likely that the weights do not increase as fast but are distributed over different areas of expertise. To adjust the speed of the CSCW increase, the CSCW build-up parameter α in equation 23 can be adjusted.

The investigation of the importance of agent's properties (research question 6) indicates that the interaction frequency is the most important feature. On the one hand this is an artifact of the simulation that only allows an increase of CSCW after an interaction took place. On the other hand this is also something observed in the social capital experiment, where the number of comments was the second most important feature for prediction of their CSCW (see table 6).

Part V

ADDITIONAL ASPECTS OF THE CSC SYSTEM

Part II introduced the concept of a contributive social capital system to create transparency and address challenges in online conversation. The basic functionality of this concept regarding the assessment of CSC with supervised learning was investigated in part III and the market system was investigated in part IV. A variety of additional investigations were performed to improve the system and examine its applicability in real life settings. These investigations are presented in the next chapters. Potential use cases of the system are discussed in chapter 16. It is also important to keep in mind that the system can potentially be misused. This crucial issue is discussed in chapter 17 at the hands of two examples. Another aspect is subjectivity, i.e. who supports whom within the social network. This may help to identify biased support from groups with a subjective agenda in mind and thereby create additional transparency. Chapter 18 presents investigations regarding this subject. Finally, chapter 19 shows how a market system could be implemented on the blockchain.

16

USE CASES FOR THE CONTRIBUTIVE SOCIAL CAPITAL SYSTEM

In part II, the concept of a system for the assessment of individual contributive social capital scores was discussed. This system is based on three pillars, extraction from online data sources, assessments via social capital markets, and individual endorsements and certifications. This chapter lists and discusses a variety of use cases of a fully implemented CSC system.

16.1 ASSESSMENT OF UNKNOWN PEOPLE IN ONLINE COMMUNICATION

One main use case of the system is to improve online communication by displaying an easily accessible score that reflects if and to what extent the user has added value to the network in the past. This score could be displayed next to a person's username in many online social networking platforms. It aggregates all available information about the user into a single value between 1 and 10 (see section 9.2.9). Based on this score, others can quickly evaluate whether they want to trust an unknown person and their message, decide whom to ask a question, or advertise their own competence.

In addition to the single score, the CSCW can be contextualized along different topics with the help of a suitable ontology (see section 9.2.8.2). This has been demonstrated in the market experiment in section 14.5.2. The division by topics should not be displayed initially to facilitate a quick assessment, but could be revealed by clicking on the overall CSC score. These topic-specific weights allow a far deeper assessment. A high score in "politics", which was one of the topics used in the social capital experiment, may add additional credibility to a post about politics from an unknown person in a social networking platform.

This CSC score can be displayed in addition or as alternative to the intrinsic reputation management systems many platforms currently employ (e.g., seller ratings on Amazon marketplace).

16.2 SEARCH FOR EXPERTS

A simple search algorithm can be implemented that allows to identify users with a certain CSCW in a defined topic. This is of special interest for three groups:

- **Private people** who are looking for help, information or a service. They can identify an individual with a high score in the respective area and ask them for help. The identified individuals are likely able and willing to assist because contributive social capital is related to expertise, knowledge, trustworthiness, and social responsibility. As discussed in section 9.3, people should have the option to hide their score or limit the availability for contact to prevent unsolicited inquiries. It is also possible to allow people to send their inquiries to an individual platform and limit the visibility to participants with high CSCW. Experts can then decide individually whether they wish to help them. This is a similar approach to what is already implemented on Q&A and threaded discussion platforms like Stack Overflow or Quora. An alternative would be to utilize a workload balancing algorithm to effectively prevent experts from being overrun with inquiries (Razzaghzadeh et al., 2017).

- **Companies** can use the CSC search for three main purposes. They can identify suitable people for hiring, identify experts on new problems within their existing staff, and identify influencers on social media and target them for marketing purposes. The providers of social networking platforms or other online expert exchanges can leverage the system on their website to improve their service and help people identify others with high CSCW.

- **Governmental institutions** can also leverage the search. They can identify knowledgeable people to involve them in the decision making in controversial topics and to sample the public's opinion. During the survey of 242 students about their experiences with social media, which is described in section 1.1, we also asked them about this possibility. 78.5% of the participants agreed that "people who have demonstrated high expertise regarding environmentally friendly and sustainable behavior should be included more in political discussions on the topic."

In all cases people should signal whether or not they want inquiries from friends or strangers to prevent them from being bothered unnecessarily. This functionality can be easily introduced into the system.

16.3 IDENTIFICATION OF BIASES AND SUBJECTIVITY IN THE PUBLIC DEBATE

A relevant part of discussions on social media is either biased or tries to provoke reactions of some kind from the readers (Paavola and Jalonen, 2015). Biased contributions are such that heavily support only one side of a debate. The topics range from simply liking or disliking a movie to much more serious events like the ongoing war in Eastern Ukraine, where a large number of misinformation and negative news is spread on social media (Paavola and Jalonen, 2015).

As transparency is an important goal of our system, it is possible to review from whom someone receives their support. This may be facilitated by leveraging blockchain technology that allows to see all transactions that occurred in the past.

In all instances with clearly defined interest groups (e.g., Democrats and Republicans in American politics), one could, therefore, visualize the flows of SCC and other metrics of support for individuals. These additional pieces of information help to better understand the motivation of an individual and to assess to what extent their messages should be trusted. An individual who is support from both parties may be more trustworthy in controversial debates than someone who previously only supported the arguments of one side. People who are close to only one side of the debate and mainly received support from this end, may be used as representatives for the group in discussions.

We investigated the identification of subjectivity and biases in five different experiments, which is described in chapter 18. Our investigations demonstrated that it is to some extent possible to visualize the degree of subjectivity not only on OSNEM but also in the scientific community.

16.4 BUSINESS MODELS FOR PRIVATE INDIVIDUALS

The CSC system allows a variety of new business models. These opportunities are mainly based on the ability to advertise one's contributive social capital with the help of the different weights, which is currently perceived as difficult by 80.6% of the 242 participants of our social media survey (see section 1.1). Therefore, these business models go hand in hand with the search algorithm mentioned in section 16.2. The variety of opportunities is exemplified at the hands of two examples:

- A student needs tutoring in the subject mathematics. By looking at the social capital weights of his fellow classmates, he can identify people who are skilled in math and those who may be willing to help because they have a high overall social capital weight. Depending on the specific implementation and the privacy settings of the participants, the student can either address the person directly or send an inquiry to the class network. People with high CSCW can choose to provide help to those inquiries or protectively advertise their skills and demand social capital currency as payment for their help. The payments further increase their CSCWs, which makes future advertising of their skills even easier.

- A freelance expert for healthcare and sports has a podcast on iTunes and a blog on which she provides information. In order to become more popular she can ask people to pay for the offered services and information in SCC instead of EUR. This increases her CSCW and thereby her visibility. Additionally, people might be more willing to pay with SCC, which they will receive again because of the basic income.

16.5 BUSINESS MODELS FOR COMPANIES

The services provided online range from news and information portals, to e-books and restaurant reviews, to name but a few. Currently these services are financed by one of the following means: advertisements, direct payments, reputation and goodwill build-up, or collecting and selling the data of the users.

The payment with SCC is a new option introduced by the system that makes the CSC build-up measurable. This has two advantages. Firstly, the CSCWs can be directly used for advertisement ("we have a high social capital in this field because we provide excellent information/services"), and secondly the earned social capital currency can be utilized by the company. The SCC can, e.g., be used to pay the writers of an online news platform, which in turn increases their CSCWs and makes it easier for them to prove their expertise on the subject.

In 2017, the accounting firm EY interviewed 1,400 Germans about their online behavior. The five most popular online services were communication (used by 81% of the participants), news and information (70%), shopping (65%), music (50%), and financial transactions (49%) (Earnest & Young, 2017). In the following, we briefly discuss how the CSC system may complement these services.

- **Online communication.** The largest online social network and thereby the largest provider of online communication is Facebook. It is the main goal of the CSC system to improve online communication and we already discussed that displaying CSCW scores on the platform may prevent the spreading of fake news and lies because it enables people to make quick assessments about the trustworthiness and competence of a source.

- **News/information.** About 44% of Americans obtain their news from social networking platforms like Facebook (Owens, 2017). For this group the same observations hold true as for online communication. A CSC system would improve their experience as it allows them to assess the trustworthiness of an article more easily. Other news publishers may use the system in the way described above, to advertise their competence and to receive and make payments with SCC.

- **Shopping.** Large online retailers can already use machine learning analysis to minimize fraudulent behavior, e.g., people ordering products without paying for them. If a person is unknown or for another reason flagged as untrustworthy, only low risk payment options, e.g. prepayment, are provided. This analysis may become obsolete if new customers can exhibit their trustworthiness with a high CSCW. The same transparency may enhance the experiences of both, sellers and buyers, on platforms like eBay or Amazon Marketplace.

- **Music.** Spotify is one of the largest providers of music online. It offers two versions, a free version with advertisements and a paid plan. Payments with SCC may be a third option. The SCC can be redistributed to the artists, who can use their CSCW increase for direct advertisements. This is especially interesting for new and lesser known musicians.

- **Financial transactions and banking.** For banks the CSCW of customers may help to assess their creditworthiness in a similar but extended way as the Schufa Holding AG (Toader et al., 2015) in Germany. Such an assessment is also the basis of the Chinese social credit system, which is discussed in section 17.2. For individuals the CSCW score and SCC payment system may also be useful for interactions with previously unknown service providers.

16.6 INTRODUCTION OF CSCW BASED VOTING MECHANISMS

Weighted voting means that the voting power of individuals is not necessarily the same but linked to other properties, which is already established in some areas of business. Common shareholders of companies are often granted the right to vote on major corporate issues at shareholder meetings (Investorpedia, 2018). Their voting power is proportional to the number of shares they hold.

The CSCWs allow something similar but instead of coupling the voting power to the monetary involvement it can be linked to a person's previous value-add to the community in different topics. This may bring advantages for individuals, companies, and governmental institutions.

1. **Individuals** can use voting based on CSCWs as recommender system for decision making in groups. By taking the weights in the relevant topics into account, the vote of experienced participants is more important. This can help, e.g., when deciding where to have dinner with a group in a new city. The people within the group who are local and/or experienced with food have a higher voting power. To what extent this improves the satisfaction of all parties needs to be investigated further.

2. **Companies** can leverage CSCW voting in a similar way, e.g., by enriching factual debates with a vote on the topic that includes experts from different departments.

3. **Governmental institutions** may also use CSCW-based voting. Successful political decisions need to be based on facts that are obtained by looking into all aspects of and points of view regarding a problem. For many issues the CSC system can help to find agreements, by first identifying the people with the highest CSCW in the respective fields (e.g., environmental protection) and then reaching a consensus with weight-based voting.

It is important to note that this possibility should not be used to undermine the democratic process our political system is built upon. It is a mere extension to facilitate decision making on individual issues that require expert knowledge. The pros and cons need to be carefully considered before each application of CSCW-weighted voting.

16.7 supporting and incentivizing altruistic behavior

The system was designed to support altruistic behavior, as described in section 5.1. People can reward others for pro-social and helpful contributions by transferring SCC into their account. This reward could also come from governmental institutions or charitable organizations. It has been demonstrated that rewards are often seen as an additional incentive to behave in a positive way (Gneezy et al., 2011). Research on threaded discussion boards also revealed that people are willing to go great lengths for virtual likes or up-votes (Richterich, 2013). This further supports the hypothesis that a SCC market may help to evoke social behavior. A SCC transfer can also be used as a simple "thank you" between friends, where a monetary transfer would be awkward.

These hypotheses are supported by 71.1% of the 242 participants of our social media usage survey (see section 1.1), who agreed with the sentence "If there was an easy way to reward people online for social activities (e.g., providing information, helping out) or for volunteerism (e.g., for the environment), I would use it."

17

POTENTIAL RISKS OF A CONTRIBUTIVE SOCIAL CAPITAL SYSTEM

17.1 SYNOPSIS

Almost everything in life has at least two sides. Even though the CSC system was designed with the best intentions in mind, it is possible that the findings may be used for purposes that are not solely for the benefit of mankind.

In recent years, there were several incidents that underline the importance of caution with regard to research on social networks and overarching reputation systems. To promote a better understanding of the potential risks of the CSC system, this section discusses two of these occurrences, the proposed social credit system of the Chinese government and the Facebook data scandal that became public in 2018. Ways to mitigate the identified risks are discussed in section 17.4.

17.2 CHINESE SOCIAL CREDIT SYSTEM

In 2008, the People's Bank of China (PBoC) announced the development of a nationwide credit rating system. The focus of this system was the assessment of the creditworthiness of an individual, similar to the Schufa Holding AG (Toader et al., 2015) in Germany. In 2014, the scope of the system was extended by the PBoC and the National Development and Reform Commission (Cheng and Shuyang, 2014). The system should additionally assess commercial activities, input from the government, and judicial information. It is this inclusion of non-economic information that evoked multiple reports in western media.

In cooperation with eight private companies, among them Alibaba and Tencent, different pilot projects were launched in 2016 in 32 Chinese cities (Ohlberg et al., 2018).

According to media reports, the government plans to make the system mandatory for all citizens by 2020.

17.2.1 *Goals of the Chinese Social Credit System*

In current official governmental plans, the goal of moral integrity appears more often than the initial assessment of creditworthiness. Following (Ohlberg et al., 2018), the main goals of the system are creating a culture of integrity, solving economic problems, and improving the governance. The goals and subgoals are listed in table 18.

Some of these goals sound like noble endeavors and others, like the reward of good behavior, are similar to the goals of the CSC system. The difference is that a central entity, i.e. the Chinese government, defines what good behavior is.

Table 18: Official goals of the Chinese social credit system (Ohlberg et al., 2018)

Create integrity	**Solve economic problems**	**Improve governance**
Restore social trust and honesty in the Chinese society	Boost market efficiency and economic growth	Enable information exchange between governmental institutions
Reward good and punish untrustworthy behavior	Prevent product piracy	Increase government credibility
Educate citizens about moral behavior	Protect customers	Protect citizen data
Educate citizens about appropriate online behavior	Improve customs services	Fight corruption

17.2.2 *Implementation of the Chinese Social Credit System*

There is not yet a social credit system that is mandatory for all citizens of China. However, there are 32 prototypes that are already in place and allow us to investigate the implementation of the system.

The system collects information about each citizen and company. Based on this data, a central score is calculated. Depending on the height of the received score, the citizens can receive rewards or even punishments.

The collected information differ from prototype to prototype. Most prototypes include credit assessments, judicial information, and input from governmental institutions. Some prototypes also record the activity on online social media, as well as what products a person bought online (R. Botsman, 2017). The latter is a direct possibility to influence a person's behavior, as buying video games may be punished but the purchase of diapers may be rewarded as they indicate that the person raises a child, which is seen as positive for society. Including interpersonal relationships in the assessment, is another method to influence the behavior of citizens. It makes it possible to punish people who associate with individuals who express criticism about the political system (R. Botsman, 2017).

There are different scales to record the social credit score. The sesame credit prototype, e.g., uses a range from 350 to 950.

There are several possible rewards for a high score. The list ranges from faster check-ins in hotels and the Beijing airport to easier loans, and more prominently displayed profiles on dating apps (R. Botsman, 2017). On the other side, people with a low score are punished. They receive lower Internet speeds, do not have the right to travel abroad, and their ability to obtain a loan or insurance is restricted. They are also banned from certain jobs, like lawyers or journalists, and face difficulty when enrolling their children into prestigious schools (R. Botsman, 2017).

17.2.3 *Risks of the Chinese Social Credit System*

There are two main causes for concern regarding the Chinese social credit system. First of all, the list of punishments is severe. It is not yet clear to what extent these punishments will be included in the final version of the system. Jeopardizing the education of one's children may, however, be enough to force people to behave in a certain way. This brings us to the second issue. The definition which kind of behavior is wanted or unwanted, i.e. which actions increase or decrease the score, is up to a central institution. Currently this entity is the Chinese government, the PBoC, or one of the eight private companies involved with the roll out of the prototypes. It is easy to imagine that the interests of the individual entities play a role when defining the criteria.

According to the English newspaper The Telegraph, the system has already banned 12 million citizens from traveling as punishment for unwanted behavior. The unwanted behavior included not paying fines or smoking in no smoking areas (TheTelegraph, 2017). These restrictions are enabled by the fact that many transactions in China are conducted via mobile pay, a smartphone-based payment option that allows a unique identification of the user. In 2016, some places did not accept cash payments anymore and the overall mobile pay volume was already at 5 trillion USD (CNBC, 2017).

A system that operates this way is the dystopia of the CSC system presented in this thesis and severely contradicts the goals defined in section 5.1.

17.3 Facebook Data Scandal

In the 2016 American presidential election, both candidates used social media to convince the public of their campaigns. This has been tried in the past but often with limited success (Hong and Nadler, 2012). Due to the progress of social networking analysis methods the effect on this election was by all appearances more significant. Cambridge Analytica, a consulting company connected to the Trump campaign, analyzed the data of at least 87 million Facebook users to improve targeted advertising for Trump's campaign (NY-Times, 2018) and claim to have had significant impact.

In the following, it is summarized what has been reported about the analysis conducted by Cambridge Analytica.

The first step was to collect a data set. Similarly to what was described in chapter 10 and chapter 12.5, for each user a variety of features was collected. The features included personal information, "like" patterns on Facebook, and excerpts from the voter register. Based on these features it was possible to predict a variety of characteristics, including the user's sexual preferences or political affiliation. (Krogerus and Grassegger, 2016)

In order to influence a person's voting preferences, Cambridge Analytica aimed at understanding each person on an individual level. This allowed them to address the individual's specific fears and motivations and was achieved with the help of psychometrics. The Facebook app "MyPersonality", which was developed at Cambridge University, allows people to assess themselves regarding to the "Big Five" personality traits (John and Srivastava, 1999). The "Big Five", also called the Ocean model, describe someone's personality according to five characteristics: openness, conscientiousness, extraversion, agreeableness, and neuroticism. Several million people used this app to assess themselves and thereby created one of the largest psychological data sets. It included the Ocean assessment as well as the "like" patterns and other Facebook features of the users. This data set was later used by the company Cambridge Analytica as ground truth. (Krogerus and Grassegger, 2016)

The Ocean model has been used in a variety of applications: to predict career success (Judge et al., 1999), subjective well-being (Sheldon et al., 1997), and differences between men and women (Schmitt et al., 2008). Cambridge Analytica used the collected data for political marketing with so-called micro targeting. With the help of personnel features collected from Facebook and the Ocean model ground truth, it was possible to predict the Ocean profiles of millions of American voters. The data set was enriched with information from other sources, like voter registers, medical information, and magazine subscriptions. Cambridge Analytica claimed to have created personality profiles of all 220 million adult Americans (Krogerus and Grassegger, 2016). Based on these profiles it was possible to create personalized advertisements that present the same message in different ways – individually tailored to the recipient. Second amendment rights that are traditionally Republican positions can for instance be pitched to undecided voters in several ways. A fearful person with a high neuroticism score may be convinced by the argument that guns are important for self defense, after they have seen the latest burglary statistics. Conservatives with high extraversion scores on the other hand may be convinced that traditional values are on the line if they are no longer able to go duck hunting with their sons, which is a family tradition. These different messages could be combined well with Trumps alternating positions on different controversial topics. Several 100 thousand personalized messages were automatically sent to users, to either convince them to vote for Trump or to abstain from voting at all. (Krogerus and Grassegger, 2016)

While it is likely that these directed advertisements had some impact (Persily, 2017), it is difficult to predict to what extent they influenced the election. It also clearly indicates that we are in a new age of marketing, where citizens have become more and more transparent. The CSC system aims to create transparency about a person's CSC along different dimensions. These dimensions reflect the person's interests and may consequently, also be used as input for directed advertisements. Additionally, the CSC system collects data for each participant, which is needed for the CSCW assessment. If this data were to be exploited by a third party, it would allow to create personality models as or even more detailed as the ones used by Cambridge Analytica. Such a misuse would lead to a major loss of trust and needs to be prevented, as discussed in the following section.

17.4 POSSIBLE WAYS TO MITIGATE RISKS IN THE CSC SYSTEM

The previous two sections detailed two controversial use cases of social reputation systems and social network analysis. The extent to which these methods are already in use, e.g. for marketing products (Golbeck et al., 2011) or movie recommendations on popular streaming services (Gomez-Uribe and Hunt, 2016), underlines their importance.

The contributive social capital system was designed as a counterweight to the efforts of large corporations, by creating transparency for everyone — not just for large companies or political parties that can afford it. Nevertheless, there is a risk that the findings are misused for goals that contradict those presented in section 5.1.

The largest potential for misuse of the system is if it were to be used for the purposes of the two discussed examples: to influence people in ways they do not want to be influenced — either to make money or to achieve political outcomes that are not necessarily in the public's best interest. To mitigate these risks and ideally prevent them from happening, one can take the following precautions.

- Setup of the CSC system in a decentralized way. One main issue of the Chinese social credit system and the Facebook data scandal is that there is only one central authority. This authority oversees all data and – with the restraint of some international laws – can use them in any way they like. The idea behind a decentralized setup is straightforward: If there is no centralized organization it cannot misuse the power vested to it, e.g., by selling the data or artificially increasing the CSCWs of people from one side of the political spectrum. A way to realize the system in a distributed way is the use of a distributed blockchain to save all user information and the market transactions.

- Another attempt may be to set the system up in a centralized way but carefully select the central entity. The reason is that a distributed setup of the CSC system does not solve all problems. The system may still be influenced from the outside and some attacks on the integrity of the system may better be countered by a single institution. An example is the hack on the distributed system "the DAO" that led to an end of the system (see section 13.2.5). The integrity of such a centralized entity is essential for the overall trustworthiness of the system. Thus, it should be a reputable organization, as discussed in section 9.2.2.

- No matter how the system is set up, most decisions should be made by the community as a whole and not by a central institution. This could be achieved by implementing a democratic voting system for all kinds of decisions, similar to the one used for the DAO (section 13.2.5).

- Finally, high transparency levels are essential to prevent misuse. If the CSC system were implemented on the blockchain, it would allow to review all previous SCC transactions. The OSNEM analysis mechanism cannot be published because it would encourage fraudulent behavior. However, regular reviews through a trusted authority may be possible and increase trust.

All of these precautions need to be implemented with regard to the design requirements discussed in chapter 9.

18

ANALYSIS OF SUBJECTIVITY REGARDING SUPPORT

The following experiments were conducted in the context of the master's theses by Valeriia Chernenko (Chernenko, 2017), Johannes Feil (Feil, 2017), and Rauf Zeynalov (Zeynalov, 2018), which were supervised by Sebastian Schams and Georg Groh in the Social Computing Research Group at Technical University Munich.

18.1 SYNOPSIS

In section 16.3, we argued that an analysis of biases and subjectivity on social media may help to understand the provenance of the support of individuals and groups.

Often, the statements on online social networking platforms that inspire the most reactions are controversial. Topics like US politics, exit from nuclear power, or global warming incite people to take sides and defend their position meticulously. In this context, *subjectivity regarding support* describes the phenomenon that proponents of certain theories will support their cause no matter what. A vivid Trump supporter, e.g., will stand behind his president's decision regarding all topics. Visualizing the flows of support a person receives from different groups can help understand the provenance of their support and thereby the position and motivation of an individual.

The transparency created by the CSC system may help to visualize such biased or subjective behavior.

To investigate means to create this visibility, we ran several experiments on different social networking platforms, a microblogging and a Q&A portal, as well as a scientometric database. These experiments have the following steps in common:

1. Identify different communities that represent groups of people with diverse opinions or characteristics with the help of clustering.

2. Visualize the flows of support of relevance for CSC between the clusters. This was achieved by using the metrics identified in chapter 12 as ground truth approximations for CSCW.

3. Compare and discuss the results based on the nature of the clusters.

In the following, we analyzed excerpts from Twitter, Facebook, a scientometrics data set, and Quora. It was possible to visualize some levels of subjectivity on a cluster-level. To extend these investigations to an individual level, we used the data that was collected for the social capital experiment (see chapters 11 and 14). With the help of the additionally collected information, we were able to also investigate CSC related subjectivity on an individual level, which is described in section 18.7.

Finally, the findings are summarized and discussed with regard to the CSC system in section 18.8.

18.2 MOTIVATION AND RESEARCH QUESTIONS

Researching subjectivity and biases enables a new level of transparency: seeing from where someone receives their support. The research questions of interest are:

1. Can users be clustered in terms of their CSC flows in a way that is meaningful regarding subjectivity in CSC assessments?
2. How can these investigations be extended to examine support on an individual level?

These questions were investigated on data sets from Twitter, Facebook, Scientometrics, Quora, and from the social capital experiment.

18.3 SUBJECTIVITY ON TWITTER

This section discusses the investigations regarding subjectivity on Twitter. These examinations follow the procedure presented in section 18.1 and are similar to the following analyses on Facebook, Quora, and the scientometrics data set. To avoid repetitions, the procedure is discussed in more detail in this section and in an abbreviated version in the following sections 18.4, 18.5, and 18.6. Whenever possible the differences between the data sets were leveraged to perform additional analyses, e.g., regarding the individual support of scientists from different fields in section 18.5.

18.3.1 *Description of Data Set, Graph, and Clustering*

The data set is the same as the one used in section 12.6. In that section, several different graphs were created based on the data set. The retweet graph, where the edges are the number of retweets, is of particular interest for analyzing contributive social capital. The reason is that retweets can be seen as symbol for appreciation. The number of retweets was also used as estimate for a CSCW ground truth value in the previous investigations in chapter 12.

The retweet graph was consequently used as basis for the analysis. It has 25,000 nodes and 311,255 edges that are weighted by the number of retweets between two people.

Different communities were identified with the help of Newman's greedy algorithm that optimizes the modularity of the resulting clustering (Clauset et al., 2004). This algorithm was selected because of its scalability for large networks. It outputs a dendrogram, which was cut at the level where the modularity of the clustering is maximum.

On Twitter, the algorithm identified 3,805 different clusters in the retweet graph. Most of these clusters are small, often consisting of only one person. Only 8 clusters contain more than 100 Twitter users and include 83% of all nodes in the network.

18.3.2 *Subjectivity Analysis*

After the identification of different communities in the retweet graph, we can investigate the mutual support between these clusters.

We introduce a subjective support matrix S that describes the contributive social capital between different clusters. For cluster C_i and cluster C_j, S can be defined in the following way:

$$S_{ij} = CSC_{C_j}(C_i) = \frac{1}{|C_i| \cdot |C_j|} \sum_{a \in C_j} \sum_{b \in C_i} CSC_a(b), \tag{46}$$

where CSC stands for contributive social capital, CSC_{C_j} is the accumulated contributive social capital of the whole cluster C_j, and $CSC_{C_j}(C_i)$ is the part of the accumulated CSC that cluster C_j received from cluster C_i. Similarly, $CSC_a(b)$ is the contributive social capital person a, who is in cluster C_j, received from person b, who belongs to cluster C_i. $|C_i|$ is the number of edges in cluster C_i, which is used for normalization as the clusters are of different size.

The metric we used as approximate ground truth for CSCW on Twitter is the number of retweets. With CSC = Number of retweets, the matrix element S_{ij} becomes the total number of retweets exchanged between both clusters, normalized by the total number of retweets in the respective clusters.

The subjective support matrix S is visualized in figure 22. The larger the value a matrix entry S_{ij} receives in equation 46, the darker is the respective square in the matrix. This visualizes the provenance of the support in relation to the respective clustering. The diagonal elements are expected to have large values, due to the relation of the clustering to the calculation of the matrix S. In the following section – and in the discussions of the related experiments on the Facebook, Quora, and Scientometrics data sets – we discuss the respective flows of support and investigate possible causes for the support. These investigations serve as proof of principle and can be directly employed to visualize support along different clusters. It is, e.g., possible to visualize the matrix S for a clustering along criteria x that is not related to CSC.

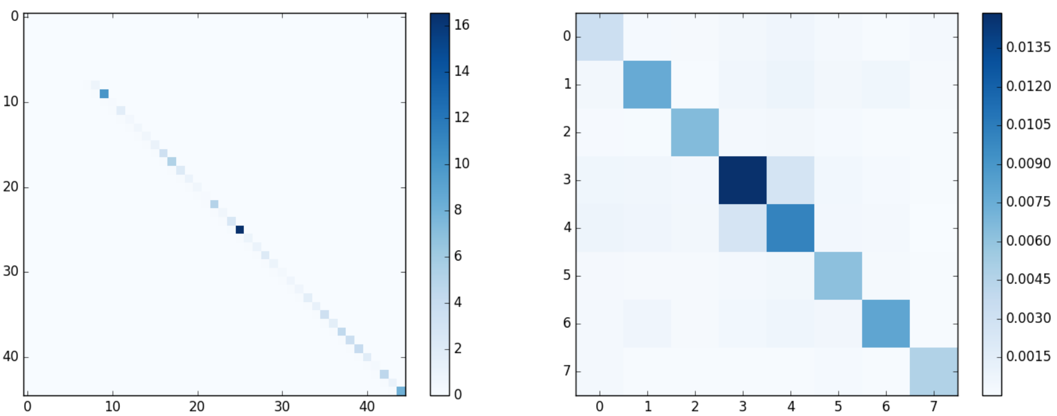

Figure 22: The subjective support matrix between the different clusters on Twitter. The left figure includes all clusters with minimum size three, the figure on the right only the largest eight clusters. The darker the square, the higher is the corresponding value in the matrix, as described by equation 46.

18.3.3 *Discussion of Results*

The visualization reflect that most of the interactions take place within the respective clusters, which can be seen in figure 22, where the diagonal values are the most prominent. This is due to the relation of the clustering and the matrix used for the visualization. The only two different clusters that exchange visible support are clusters 3 and 4. The reason for this support may be due to the topical overlap between both clusters, which is primarily UK politics and news. The topical focus of each cluster is listed in table 19. The descriptions were created by reviewing the 50 users with the highest CSC (number of retweets) of each cluster. It is apparent that most of the users are from the UK and that many of the discussed topics revolve around news, politics, football, entertainment, and religion. Within the clusters – at least among the top 50 users – there is a large topical overlap. In clusters 1, 3, and 4 the overlap was 90% or larger. Clusters 1, 3, and 4 all deal with UK politics. Cluster 1 includes many governmental institutions, while clusters 3 and 4 include accounts mainly associated with news and journalism. To better understand these clusters, we further investigated the accounts associated with them. We detected six accounts directly related to the Conservative Party and four accounts directly related to the Labour Party. Five of the six Conservative Party accounts belong to cluster 3, while all four Labour Party accounts belong to cluster 4. It is reasonable to assume that people of different political orientation belong to different social capital clusters, which may be the case with clusters 3 and 4. However, this hypothesis is based on little observed data and can consequently only be used as indication and not as proof. The retweets between the clusters that are visualized in figure 22 may be a sign that the accounts that got retweeted, represent positions that both sides of the political spectrum can agree on.

18.4 SUBJECTIVITY ON FACEBOOK

18.4.1 *Description of Data Set, Graph, and Clustering*

The data set used for the investigation of subjectivity on Facebook is the same we used for the contributive social capital investigations in section 12.5. Thus, we can use the conclusions of that section as input for the subjectivity analysis. As an alternative ground truth value on Facebook, we identified the number of likes. For the investigations of CSC flows and support between clusters, we consequently use the like graph. In this graph there are 11,629 nodes connected by 137,951 edges that are weighted by the number of likes exchanged between two people.

Newman's modularity-based greedy algorithm identified 240 clusters on the like graph. The majority of these clusters are small, often including only one node. There are 24 clusters with more than 100 nodes and only eight clusters that include over 500 users.

18.4.2 *Subjectivity Analysis*

With CSC = Number of likes, we can use equation 46 to analyze the CSC-relevant interactions between the clusters. The resulting matrix is provided in figure 23.

Cluster	Topical focus	Top 5 users
Cluster 0	International news and journalism (66%)	*BBC Breaking* (1552 retweets), *BBC World* (807), *The Economist* (742), *Al Jazeera English* (733), *Mashable* (independent news website, 677)
Cluster 1	UK politics, government and organizations (80%)	*David Cameron's office* (465), *Joseph Rowntree Foundation* (social policy research charity, 355), *UK public and voluntary sectors from the Guardian* (310), *UK Foreign and Commonwealth Office* (234), *UK Parliament* (227)
Cluster 2	Sports and football (mostly about FC Liverpool, 90%)	*FC Liverpool supporting account* (1105), Piers Morgan (presenter of the TV show Good Morning Britain, 904), *Henry Winter* (Daily Telegraph Football Correspondent, 830), *Brian Durand* (FC Liverpool supporter, 825), *Official FC Liverpool account* (761)
Cluster 3	UK politics, news and journalism (92%)	*The Queen* (1512), *Sky News Break* (1002), *Tim Montgomerie* (editor and columnist, 907), *Andrew Neil* (BBC presenter, 764), *House of Twits* (website about UK politics, 625)
Cluster 4	UK politics, news and journalism (94%)%	*Paul Waugh* (Executive Editor of Huffington Post UK, 1110), *Guardian News* (747), *Owen Jones* (Independent columnist, 707), *BBC Radio 4* (676), *Jon Snow* (presenter of Channel 4 News, 656)
Cluster 5	Entertainment (70%)	*Stephen Fry* (British actor and writer, 806), Hungton Post UK (331), *Londonist* (Events in London, 154), *Phillip Schoeld* (TV presenter, 138), *ITV Football* (129)
Cluster 6	Birmingham and West Midlands (70%)	*West Midlands Police* (352), *Birmingham Mail* (255), *Birmingham Post* (240), Simon Pegg (actor and comedian, 170), *Claire Spencer* (Councilor in South Birmingham, 167)
Cluster 7	Religion and spirituality (82 %)	Rihanna (287), *Proverbs and Quotes* (215), *Founder of ProductiveMuslim.com* (199), *Astrology account* (188), *Inspirational Oprah Fan site* (178)

Table 19: Topical focus and top 5 users of the eight largest clusters identified in the Twitter data set. The percentage in brackets states how many of the top 50 users in the cluster fit the description (topic focus) of the cluster. Users in the top 5 who fit the description are written in italics.

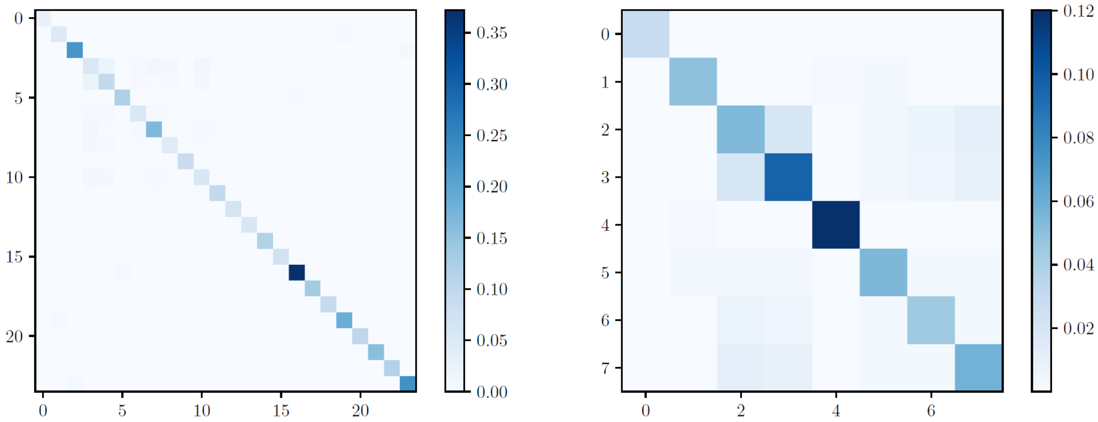

Figure 23: The subjective support matrix between the different clusters on Facebook. The left figure includes all clusters with minimum size of 100, the figure on the right only the largest eight clusters. The darker the square, the higher is the corresponding value in the matrix, as described by equation 46.

18.4.3 *Discussion of Results*

Similar to Twitter, we observe that most interactions that are of relevance for contributive social capital take place within the clusters — the diagonal elements of the matrix have the by far largest values.

There is also some exchange of likes between different clusters. Especially clusters 2 and 3 are active; with each other and with clusters 5, 6, and 7.

The clusters were formed artificially with the help of global graph clustering. We could not identify common topics that distinguish the different clusters from each other. This is mainly due to the nature of Facebook, which – at the time of the crawl – focused on everyday conversations between people and included less topic-specific discussions or news sharing as, e.g., Twitter.

We could, however, identify commonalities regarding the place of residency of the people in the different clusters, which allows us to investigate which regions particularly support each other. The top places of residency are listed in table 20.

It is apparent that the majority of users live in the United States. The different places of residency may also be able to explain the inter-cluster support that we have seen. Most people from clusters 2, 3, and 7 are from Northbridge, Massachusetts. Clusters 5 and 6 who also interacted with each other and the previous three clusters, also include predominantly cities in Massachusetts. The extraordinarily large intra-cluster support of cluster 4, may be due to the fact that over 40% of people in this cluster live in three cities that are less than 40 miles from each other.

While these geographic investigations may be able to explain part of the subjective support that we have observed, one does not need to stop there. A variety of publications investigate to what extent user characteristics can be extracted from Facebook, which is similar to our investigations in chapter 12.5. David et al., e.g., trained a classifier to predict the political orientation of Facebook users with high accuracy (David et al., 2016). With the help of such a classifier, one could identify different political clusters and investigate the CSC flows between them to identify users who

Cluster	Cluster Size	Top places of residency
Cluster 0	797	Chicago (126); Wheaton (53); Los Angeles (51); St. Louis (35); Brisbane (15)
Cluster 1	1199	Wenham, Massachusetts (92); Beverly, Massachusetts (72); Boston (71); New York City (22); Los Angeles (21); Salem, Massachusetts (17); Gloucester, Massachusetts (17); Concord, New Hampshire (12); Washington, D.C. (10)
Cluster 2	1515	Northbridge, Massachusetts (182); Worcester, Massachusetts (126); Whitinsville, Massachusetts (84); Uxbridge, Massachusetts (61); Douglas, Massachusetts (54); Boston (41), Grafton, Massachusetts (34); Milford, Massachusetts (25); Millbury, Massachusetts (24)
Cluster 3	649	Northbridge, Massachusetts (112); Whitinsville, Massachusetts (34); Uxbridge, Massachusetts (29); Boston (29); Worcester, Massachusetts (18)
Cluster 4	638	Gloucester, Massachusetts (189); Boston (69); Rockport (21)
Cluster 5	537	Boston (45); Holliston, Massachusetts (32); Whitinsville, Massachusetts (28); Worcester, Massachusetts (20); Sutton, Massachusetts (10)
Cluster 6	545	Worcester, Massachusetts (83); Sutton, Massachusetts (31); Boston (18); Westfield, Massachusetts (14); Millbury, Massachusetts (11); Amherst, Massachusetts (10)
Cluster 7	657	Northbridge, Massachusetts (54); Whitinsville, Massachusetts (35); Uxbridge, Massachusetts (33); Boston (31); Worcester, Massachusetts (30); Douglas, Massachusetts (16)

Table 20: The locations stated as place of residency of at least 10 people from each of the eight largest clusters identified on the Facebook data set. For lesser known cities the state is listed as well.

are supported from all sides of the political spectrum. The findings of this section can directly be leveraged for these investigations.

18.5 SUBJECTIVITY IN SCIENTOMETRICS

18.5.1 *Description of Data Set, Graph, and Clustering*

The analysis of subjectivity in scientometrics allows to investigate from which communities scientists, or group of scientists, receive the most citations. It was performed on the ArminMinder data set described in section 12.7.

The number of citations, a universal symbol for appreciation in the world of science, was used in that chapter as ground truth alternative for CSC. As corresponding graph for the analysis we, therefore, use the citation graph between all authors in the data set. This graph has 99,178 nodes and 2,331,154 edges.

The clustering of this graph was also performed with Newman's greedy algorithm which optimizes the modularity of the resulting clustering. It identified 4,646 communities, which mainly consist of only 1 author. Similar to Twitter, only eight clusters have more than 100 nodes. 95% of the authors belong to one of the eight large clusters.

18.5.2 *Subjectivity Analysis*

With the help of equation 46 and CSC = Number of citations, we can again calculate the subjective support matrix S. This matrix displays the normalized number of citations between the different clusters, and thereby the origin of the authors' CSCWs. The matrix is visualized in figure 24.

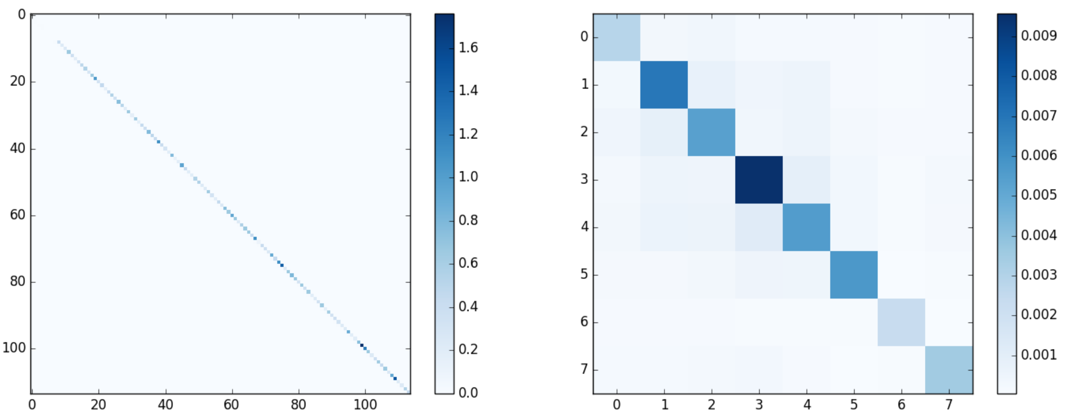

Figure 24: The subjective support matrix between the different clusters in the ArminMiner data set. The left figure includes all clusters with minimum size of 3, the figure on the right only the largest eight clusters. The darker the square, the higher is the corresponding value in the matrix, as described by equation 46.

18.5.3 *Discussion of Results*

Similar to what we have seen in the previous two experiments, the main area of support in the form of citations comes from the community the author belongs to.

With the help of the LDA topic distribution of the ArminMiner data set, we computed the average topic distribution $\overline{\theta}_{C_i}$ for each cluster C_i. In table 21 these topic distributions are illustrated by listing the five most active topics in each cluster and the topics that are more active in the cluster than in any other one.

There are some topics that are present in most of the clusters, which is due to the thematic overlap of the crawled data set. These topics, e.g., "statistics" describe methods that are relevant in most mathematical and computer science related research areas. The topics of high relative activity (column 3 in table 21) are better suited to describe the clusters. It is apparent that the scientists in one cluster publish in related fields and that there are distinctions between the clusters. While cluster 2 focuses on machine learning, cluster 4 mainly includes scientists that publish in mobile computing and distributed systems.

The flow of CSC between the clusters is visualized in figure 24. The figure on the right displays the eight largest clusters. Support between the different clusters is hardly visible. The highest inter-cluster support is between cluster 3 and cluster 4. Cluster 3 focuses on hardware solutions and software engineering, which are areas of importance for scientists publishing about mobile computing and distributed systems (cluster 4). The flow of social capital, therefore, makes sense, as scientists in cluster 4 appreciate the hardware-related work of scientists in cluster 4 that made their work possible in the first place.

The same analysis was performed with soft clustering instead of hard clustering, i.e. each scientist could be attributed to several clusters. This produced similar results but with an expected larger inter-cluster support.

To extend this research and investigate the social capital flow on an individual level, we introduce the subjectivity ratio $r_s(a)$. $r_s(a)$ is a measure for how much of an individual scientist a's CSC comes from within their own community:

$$r_s(a) = \frac{\sum_{b \in C_i} CSC_a(b)}{\sum_{d \in \sum_j C_j} CSC_a(d)}, \tag{47}$$

where $\sum_{b \in C_i} CSC_a(b)$ is the sum of all contributive social capital (in this case the number of citations) that person a received from all other authors in their own cluster C_i.

The calculated ratio is visualized in figure 25. For authors with low citation count the subjectivity ratio $r_s(a)$ is often binary, especially if they only received one citation from either within or from outside their community. This becomes apparent by looking at the histogram of the full data set (blue) or the top half (green). For scientists with larger citation counts (red, cyan, and violet in the histogram), the distribution is uniform between 0 and 0.8, with a mean value around 0.43. This reveals that there are almost no authors with high citation counts who only received support from within their clusters. While this balanced provenance of support may be expected in the open minded scientific community, we cannot assume that the same is the case in other communities. On social media one can observe the echo chamber phenomenon:

Cluster	Most active topics	Relatively active topics
Cluster 0	Business Informatics and Economics (13.6%), Social Computing (13.0%), Robotics (9.0%), Education and Human-Machine Interfaces (7.9%), Statistics (6.4%)	Robotics, Computer Graphics and Computer Vision, Multimedia and Games, Education and Human-Machine Interfaces
Cluster 1	Algorithms and Complexity Theory (14.0%), Linear Algebra and Optimization (7.4%), Graph Theory (6.9%), Statistics (5.6%), Artificial Intelligence (5.3%)	Algorithms, Complexity Theory, Graph Theory, Scheduling, IT Security and Testing
Cluster 2	Statistics (9.8%), Social Computing (8.5%), Natural Language Processing (8.4%), Web Search and Text Mining (8.2%), Artificial Intelligence (7.8%)	Machine Learning, Natural Language Processing, Artificial Intelligence, Databases, Web Search and Text Mining
Cluster 3	Parallel Computing and Memory Management (16.4%), Hardware and Circuits (12.0%), Software Engineering (11.3%), Theoretical Informatics (9.9%), Statistics (4.3%)	Hardware and Circuits, Parallel Computing and Memory Management, Software Engineering, Theoretical Informatics
Cluster 4	Networks (13.6%), Internet and Web Services (8.1%), Social Computing (6.9%), Statistics (6.4%), Mobile Computing (6.2%)	Internet and Web Services, Networks, Mobile Computing, Distributed Systems
Cluster 5	Numerics (12.7%), Linear Algebra and Optimization (12.6%), Social Computing (6.7%), Software Engineering (6.0%), Statistics (5.6%)	Numerics, Linear Algebra and Optimization
Cluster 6	Statistics (11.7%), Linear Algebra and Optimization (11.3%), Medicine (8.1%), Robotics (7.0%), Artificial Intelligence (6.0%)	Statistics, Medicine
Cluster 7	Social Computing (18.9%), Business Informatics and Economics (11.9%), Software Engineering (9.6%), Statistics (8.0%), Artificial Intelligence (5.6%)	Business Informatics and Economics, Social Computing

Table 21: The five most active topics in each of the eight main clusters identified in the academic network and the topics that are more active than on any other cluster (relatively active topics).

people desire to reinforce their opinions, which leads to a limited exposure regarding balanced online political information (Garrett, 2009). This likely leads to people whose support mainly originates from within their own clusters.

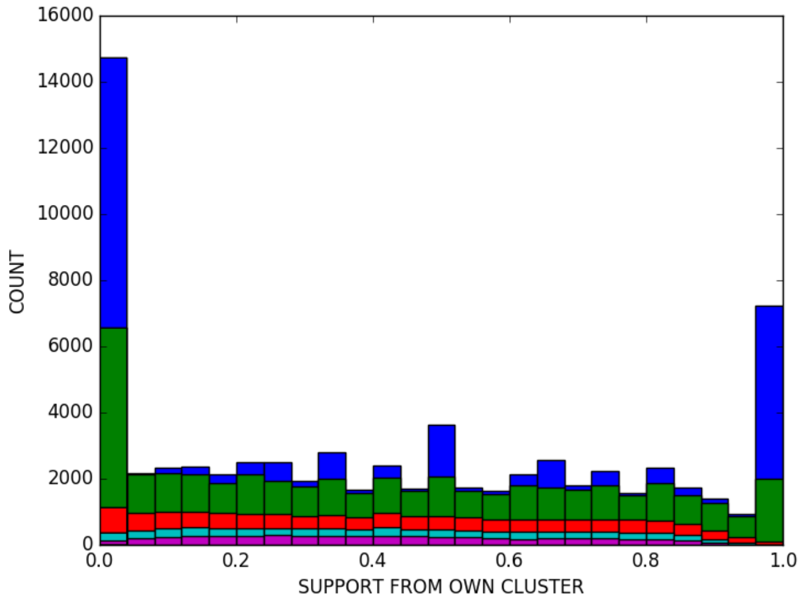

Figure 25: The subjectivity ratio r_s for different groups of authors. A high subjectivity ratio indicates that the author receives most of their support from within their own community, lower values indicate support from many different sides. Visualized are different groups of authors sorted by their CSCW: Full data set (blue), top 50% of CSCW (green), top 20% (red), top 10% (cyan), and top 5% (violet).

The subjectivity ratio $r_s(a)$ is only one exemplary metric to investigate the subjectivity of support on an individual level. Of similar interest may be the opposite measure of how much support does come from outside of the own community, or how diverse the support is. A measure for the latter could be the ratio between the number of clusters from whom person a receives support and the total number of clusters a interacted with. Further investigations of individual subjectivity, especially with regard to only two clusters, are described in section 18.7.

18.6 SUBJECTIVITY ON QUORA

18.6.1 *Description of Data Set, Graph, and Clustering*

As question and answer portal, Quora enables discussions in many different topics. By visualizing the subjectivity of support on Quora we want to investigate how much support people receive from within their community and from other clusters.

The first step is to identify clusters. For the analysis we used the same data set that was described in section 12.8. As there are no natural clusters on Quora that we can leverage, we used graph clustering to identify the different communities. The number of up-votes a person received was used as approximate ground truth for CSCW in section 12.8. The up-vote graph consequently represents the flow of CSC in the Quora data set. It has 3,069 nodes and 150,766 edges that are weighted by the number of up-votes between the different users.

Newman's modularity-based greedy algorithm was used for the clustering of the up-vote graph. It identified 9 clusters, one of which only includes 3 users. For the following analysis we will, therefore, focus on the largest 8 clusters.

18.6.2 *Subjectivity Analysis*

To analyze the support between these clusters, we can again use equation 46 and calculate the subjective support matrix. As CSC value the number of up-votes is used. The matrix is visualized in figure 26.

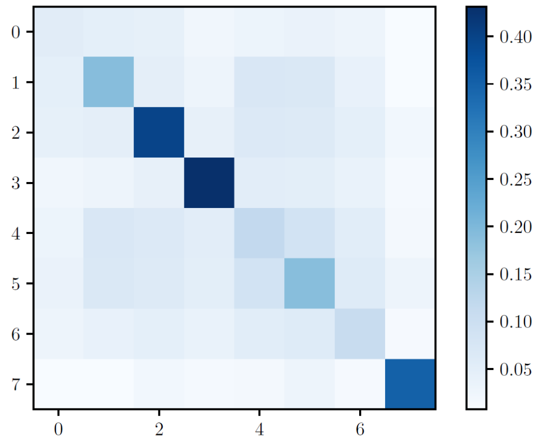

Figure 26: The subjective support matrix between the eight largest clusters on Quora. The darker the square, the higher is the corresponding value in the matrix, as described by equation 46.

18.6.3 *Discussion of Results*

On Quora we observed more inter-cluster activity compared to the other networks. Clusters 0 and 4 only displayed slightly more intra-cluster than inter-cluster up-votes. Additionally, much of the support of people attributed to clusters 1, 4, and 5 came from the respective other clusters.

As Quora fosters the exchange of ideas in different topics, we investigated the most discussed topics of each cluster to identify reasons for these exchanges. These topics are listed in table 22.

It is clear that the clusters cannot be categorized based on the discussed topics alone. There is a large overlap of topics between the clusters; history appears in seven of the eight clusters, science and politics in six. Each cluster includes a variety of topics from different domains.

Nevertheless, we observe an especially large overlap between the clusters whose people particularly support each other with up-votes. Clusters 4 and 5 have five of

Cluster	Cluster Size	Top 10 discussed topics
Cluster 0	664	History (317), Politics (224), Technology (207), Science (203), Economics (173), The United States of America (158), World History (154), Military History and Wars (154), Books (153), Movies (143)
Cluster 1	301	History (138), The United Kingdom (113), Science (106), Technology (102), Politics (99), Books (89), Europe (72), Visiting and Travel (71), Economics (68), Movies (67)
Cluster 2	101	History (138), The United Kingdom (113), Science (106), Technology (102), Politics (99), Books (89), Europe (72), Visiting and Travel (71), Economics (68), Movies (67)
Cluster 3	77	History (26), Religion (26), Christianity (25), Science (25), Books (25), Writing (21), Philosophy (19), Technology (19), Psychology (19), Politics (17)
Cluster 4	663	Science (264), History (207), Technology (198), Politics (173), Philosophy (157), Books (156), Politics of the United States of America (156), Movies (137), Writing (133), Psychology (131)
Cluster 5	578	Books (243), Writing (214), Quora (200), Food (200), Psychology (193), Music (188), Movies (187), History (186), Life and Living (180), Human Behavior (165)
Cluster 6	264	Technology (57), Politics (55), Science (55), Economics (52), Politics of the United States of America (48), History (46), Books (45), The United States of America (42), Quora (42), Philosophy (42)
Cluster 7	35	Writing (19), Life and Living (17), Psychology of Everyday Life (16), Education (15), Movies (15), Philosophy of Everyday Life (15), India (14), Books (14), Quora (14), Creative Writing (14)

Table 22: Top 10 topics of the eight clusters identified on the Quora data set. The value in brackets states the number of questions discussed in the topic.

the ten topics in common (History, Books, Movies, Writing, Psychology), 1 and 4 six topics (History, Science, Technology, Politics, Books, Movies), and 1 and 5 three topics (History, Books, Movies). Among their top ten topics, clusters 0 and 4, include many of the overall most common topics (e.g., History, Science, Politics). This may explain why people belonging to these clusters receive up-votes from all other clusters. The people in cluster 3 mainly support each other and receive relatively little up-votes from other clusters. This may also be explained by the discussed topics. Religion and Christianity, the second and third most popular topics of cluster 3, are not present in the top ten of any of the other seven clusters.

Of particular interest for the social capital system is, if there are two clusters that have supported different positions in the past. An individual who has received equal support from both clusters may be an ideal candidate to moderate a debate or to fact-check the arguments of both sides. This can be investigated on an appropriate data set with the methods discussed in this chapter.

18.7 SUBJECTIVITY IN THE SOCIAL CAPITAL EXPERIMENT

The social capital experiment consisted of two main investigations: analyzing a social network, as described in chapter 11, and examining transactions in a social capital market, which is discussed in chapter 14. The data collected in this experiment, especially with regard to the social networking platform, can also be investigated regarding the subjectivity of support.

18.7.1 *Description of Data Set and Clustering*

The advantage of this data set is that the participants provided additional details about themselves. The gender, age, and country of origin are known with a higher accuracy than on the other OSNEM platforms.

We use this data to identify natural clusters between the users and investigate the support an individual receives from the cluster they belong to and from the respective other cluster.

These investigations are exemplary for all investigations with two clusters and can, therefore, easily be extended to cover the problem at hand.

Of the 242 students in the data set, 204 provided a full set of information. The following sections describe the analyses regarding two different groupings. These investigations were conducted with all 204 students and not only with the 165 participants for whom a ground truth assessment by someone else was available. Thus, the percentage values stated in the following do not exactly coincide with the numbers in chapter 11.

18.7.2 *Clustering of the Students by Gender*

There are 166 (81%) male and 38 (19%) female students. On average the female students were more active on the social networking platform. Even though male students accounted for 81% of the participants, they were only responsible for 68% of the likes.

User ID	Gender	Female likes	Male likes	Male likes share	P-value
633	m	1	14	93.3%	0.048
54	m	2	23	92.0%	0.009
1059	m	3	22	88.0%	0.050
131	m	16	63	79.8%	0.039
868	f	18	18	50.0%	0.020
49	m	10	6	37.5%	0.013
1375	f	11	6	35.3%	0.007
2287	m	6	3	33.3%	0.032
283	m	6	3	33.3%	0.032
974	m	6	3	33.3%	0.032
1054	m	9	4	30.8%	0.006
535	f	4	1	20.0%	0.037
63	f	3	0	0.0%	0.031

Table 23: Statistics of the 13 students in the social capital experiment data set for whom the null hypothesis was rejected. The null hypothesis states that a person receives 68% of their likes from male students.

To investigate whether an individual received extraordinarily large support from the own or the other gender, we performed a binomial test. The null hypothesis of this test is that 68% of the likes the participant received came from male students.

There were 13 students for whom the null hypothesis was rejected at a significance level of 5%, i.e. the p-value for the binomial test was below 0.05. All 13 students are listed in table 23. Three of the four women for whom the null hypothesis was rejected received significantly more likes from other women. The fourth woman received 18 likes from both, men and women. There were four men who received significantly more likes from other men and had a ratio of about 80% male to female likes. The other five men only received about a third of their likes from other men, which indicates that their contributions were especially appreciated by women.

By themselves these values are nondescript. If there is a controversial debate, however, the voice of someone who demonstrated balanced viewpoints in the past, might have a different weight than someone who mainly supported one of the groups in the past.

We investigated the data set manually and could not identify a debate about a topic for which the gender might have been relevant. However, the four male students who received primarily male likes did indeed engage in conversations that might be considered "typically male".

User ID	Country	German likes	Int. likes	German likes share	P-value
189	Germany	10	0	100.0%	0.0171
1893	Germany	12	1	92.3%	0.0226
69	Germany	15	2	88.2%	0.0247
67	Germany	27	6	81.8%	0.019
131	Germany	58	21	73.4%	0.0368
344	Germany	69	25	73.4%	0.0252
1733	India	5	10	33.3%	0.0316
1318	Egypt	6	13	31.6%	0.0088
2287	Germany	2	7	22.2%	0.0322
1715	Germany	2	8	20.0%	0.0089
1448	India	1	7	12.5%	0.0063
1845	Germany	1	7	12.5%	0.0063
974	Germany	0	9	0.0%	0.0002

Table 24: Statistics of the 13 students in the social capital experiment data set for whom the null hypothesis was rejected. The null hypothesis states that a person receives 62% of their likes from German students.

18.7.3 *Clustering of the Students by Country of Origin*

In a similar fashion we can investigate whether the like patterns depend on the country of origin. There were students from 34 different countries in the data set with the majority being from Germany (65%). For the subjectivity investigations, the students were grouped in two clusters: international (72) and German students (132). The international students were slightly more active on the platform, as Germans only provided 62% of the likes. The null hypothesis was consequently that someone received 62% of their likes from German students.

This hypothesis was rejected for 13 students, who are listed in table 24.

Investigating this table does not reveal any systematic bias between German and international students. There are six German students who received about three quarters or more of their support from fellow Germans. On the other end of the spectrum, there are also two German students who received the majority of their support from international students. For only three international students the null hypothesis was rejected. These three students are from India and Egypt and received only little support from German students (a third or less of their likes). These irregularities may indicate that the respective students discussed topics that are of particular interest to international students (e.g., Visa formalities), or that are only of relevance for Germans. Upon manual investigation of the posts and comments of the 13 students, both hypotheses have to be rejected. There are no topics present that directly cater to one of the two groups. It becomes apparent, however, that writing style and topic selection varies significantly from person to person.

18.7.4 *Discussion of Subjectivity in the Social Capital Experiment*

As an example for subjective support on an individual level, the data set was divided into two groups. One division was according to gender, the other about nationality. The investigated metric was the number of likes, which is the feature with the closest correlation to contributive social capital (see table 6). For each participant of the experiment, we tested the null hypothesis that they receive likes in a way that represents the overall distribution of likes between the different groups.

For both groupings, the null hypothesis was rejected for 13 students. This also means that over 93% of students received likes in a balanced way, with neither too much or too little support from any gender or nation.

While it was not possible to exactly pinpoint the reason for these discrepancies, the findings can be utilized to indicate potential biases in future discussions that require a balanced viewpoint. The same analysis methods can also be directly used for any other data set.

18.8 Summary of Results

This chapter presented investigations regarding subjectivity on different OSNEM platforms and a scientometrics data set. In this section, we summarize the results following the research questions from section 18.2.

The first question was, whether the users can be clustered in terms of their CSC flows in a way that is meaningful regarding subjectivity in CSC assessments. We have used two methods to identify different communities. Firstly, we used graph clustering of a suitable graph that was extracted from the respective data set. And secondly, we used attribute-based clustering by using information about each user to assign them to a group. The communities identified with graph clustering could be further classified with the help of topic modeling or manual investigation. With this help we were able to identify different clusters based on gender, country of origin, thematic interests, or political orientation. These groupings can be extended to investigate subjectivity and biased support in all kinds of different communities.

Once these clusters were identified, the differences between inter- and intra-cluster support could be investigated. The experiments demonstrated that there were significant variations. Most of the contributive social capital (in the experiments represented via retweets, likes, citations, and up-votes) of a person comes from within their own cluster, which is due to the nature of the clustering. The social capital flows can be visualized with the help of the subjective support matrix, whose diagonal elements were significantly larger than the other elements. This matrix allows to directly assess how clusters interact with one another. The reason for the preferred support from the own community may be homophily, a principle introduced in 1954 (Lazarsfeld et al., 1954). It describes the tendency of people to associate and bond with others who are similar and behave in a similar way. In online social networks homophily is often described as shared group interests or as group mindset (Brown et al., 2007).

We also identified some clusters that collaborate more than others. This could be attributed to either similar characteristics, e.g., residency in the same region, or to similar thematic interests, like science and technology. The most inter-cluster support of all investigated data sources could be found on Quora. This may be due to the nature of Quora, which is a Q&A platform that promotes the open discussion of questions in all kinds of topics, where friendships are less explicit (Gómez et al., 2008). It also supports the conclusions of chapter 12 that communities like Quora and Stack Overflow are of particular importance for the analysis of CSC.

The second research question asked about the possibility to extend these findings to visualize CSC related subjectivity on an individual level. This was investigated on the scientometrics data set and the social capital experiment data. The subjectivity ratio $r_s(a)$ of an author a displays what percentage of their citations comes from within their own cluster. Especially for authors with high citation numbers it is increasingly unlikely that they were only cited from outside their own cluster. On the social capital experiment data, we demonstrated that it is possible to visualize the support ratios between different groups. At the hands of two examples, we demonstrated that it is possible to identify users who received the majority of their up-votes from only one of the two groups.

The experiments also revealed differences between the data sources. The CSC clusters on Facebook, e.g., were mainly characterized by relationships and geographical proximity rather than topical interests, as observed on Twitter and Quora. This may indicate that the latter two platforms are more suited for the investigation of biases and subjectivity regarding specific topics. The most inter-cluster support was observed on the Q&A portal Quora, which indicates a higher mixing of the respective clusters than on the other platforms. The relatively little inter-cluster support in the citation network may be explained by the fact that most scientists primarily cite others whose work is related to their own research interest. Nevertheless, the subjectivity investigations on the individual level indicate that the majority of scientists with many citations receive citations from outside their own cluster as well.

To summarize, it is to some extent possible to visualize subjective support on a cluster as well as on an individual level, which was the content of the tenth overarching research question from section 1.2. These findings can directly be employed to improve the CSC system, e.g., by visualizing the subjective support matrix or the subjectivity ratio $r_s(a)$ for all participants and platforms. The clustering that is the basis for these investigations can be adjusted to address different problems, e.g., the party affiliation in political debates.

Additionally, it might be interesting to visualize specific aspects of subjectivity for some individuals. Unbiased political newspapers might, e.g., be identified by looking at the support they receive from the different parties — similar to our findings about the Labour and Conservative Party in Great Britain in section 18.3.3.

19

IMPLEMENTING A MARKET SYSTEM ON THE BLOCKCHAIN

19.1 SYNOPSIS

One of the main challenges of the CSC system is how to implement it in a safe, transparent, and easy-to-use manner. One way to achieve that is to leverage blockchain technology. In this chapter we briefly review the available technology and describe how the market system was implemented as a smart contract on Ethereum as a proof of concept.

19.2 MOTIVATION AND RESEARCH QUESTIONS

A fundamental part of the CSC system is the social capital market, which is described in section 7. During the social capital experiment, we used a market that implemented the following features:

- Users can register with unique profiles.
- Users can receive and spend social capital currency.
- Based on these transactions, a CSCW is formed for every participant

On the limited scale of the experiment, this implementation worked well. However, if it were to be implemented on a larger scale it has a major disadvantage. This disadvantage is security. Security is an essential part of the CSC system, which needs to be fair, transparent, and to be trusted by the users in order to achieve its goals. The version we used for the experiment does not fulfill these high standards. All transactions are saved in a single database, which may be hacked or altered by the administrator, without anyone knowing. This single point of attack has been addressed with new technology that is based on the blockchain (Tapscott and Tapscott, 2016). In this context there are two research questions:

- Is there a technology that can be used to directly implement a social capital market following the principles introduced in chapter 7 and that is transparent, safe, and easy-to-use?
- What are potentials and limitations when implementing the CSC market system as a smart contract on the Ethereum blockchain?

19.3 SUITABLE TECHNOLOGIES: CRYPTOCURRENCIES AND THE BLOCKCHAIN

In recent years, the popularity of cryptocurrencies like Bitcoin or Ethereum rapidly increased. In May of 2018 the market capitalization of these two currencies alone was 124 respectively 58 billion Euros[1]. Cryptocurrencies provide the possibility for

1 https://coinmarketcap.com/de/ (retrieved: 2018-05-11)

decentralized, instant, and borderless payments. The main idea behind it is that they are not saved physically or in a file, but as transactions within a blockchain. In the following, we describe its functionality based on (Tapscott and Tapscott, 2016).

The blockchain is distributed and runs on the computers of volunteers around the world. There is no central database, which could be hacked. The blockchain is public and can be reviewed at any time. It is encrypted with public and private keys. In a defined frequency – every 10 minutes for Bitcoin – all new transactions are verified and saved in a new block, which is connected to the preceding blocks, thereby forming a chain. As every transaction is time-stamped and saved in the distributed blockchain, it becomes nearly impossible to steal a Bitcoin without the owner's private key. To do that, the thief would need to alter the history of the coin, which would be visible to every participant and thereby the theft could be detected.

The cryptocurrency Ethereum goes beyond pure transactions of currency. It allows the implementation of smart contracts that enable two or more parties that do not necessarily trust each other, to conduct transactions without a trusted third party. As summarized by Kosba et al. "in the event of contractual breaches or aborts, the decentralized blockchain ensures that honest parties obtain commensurate compensation" (Kosba et al., 2016). These smart contracts can be programmed freely and then implemented on the blockchain, where interested parties can interact with each other based on them. Smart contracts can also be used to implement the functionality of the CSC system's market.

The decentralized setup, the transparency, and security are in line with the requirements of the market system (research question 1). In the following we describe the experimental implementation of the CSC system as a smart contract and then draw conclusions about the potentials and limitations of the blockchain for the CSC market system to investigate the functionality.

19.4 IMPLEMENTATION OF THE CSC MARKET AS A SMART CONTRACT

Ethereum has its own programming language to write the smart contracts. This language is called "Solidity" and is similar to Javascript (Dannen, 2017).

The social capital market was implemented as smart contract on the Ethereum blockchain. The following functionalities were implemented as rules in the smart contract.

- **Universal basic income**: Each account that is registered as a participant receives a universal basic income in SCC. After a fixed time period – in this case every month – a defined amount of social capital currency is transferred to each participant. This money can be used freely by the users.

- **SCC transactions**: As Ethereum is mainly used as currency, it is not surprising that SCC transactions can easily be implemented as part of the smart contract CSC system. The transactions work as described by equations 20 and 21, but even more complex systems are possible, e.g., with multiplicators that increase the received amount of SCC based on the sender's CSCW. It is important to implement the full functionality directly at the beginning as later alterations are not possible.

- **Build-up of CSCW**: This is the part of the CSC system that distinguishes it the most from other markets. Every transaction of CSC is supposed to increase the recipient's CSCW.

  The smart contract contains a set of rules which are executed depending on the actions of the users. The rules directly reflect the CSCW build-up mechanism described in chapter 7. If SCC is, e.g., transferred from user A to user B, B's weight is automatically increased based on equation 27. A distinction between topics can be implemented in the same way, e.g., with the help of differently flavored currencies. A requirement of the Ethereum smart contracts is that they behave exactly the same way for every user. This is also a requirement for the CSC system whose basic functionalities are supposed to be the same for all users.

- **Time-value of SCC**: The time value of SCC has not been described in the previous chapters. The idea behind it is that unused SCC will be destroyed automatically after a certain time period. The advantage is that it would likely reduce inflation that we have experienced in the market experiment (chapter 14.5.3) and the simulation (chapter 15.5). It is also likely that this time value encourages active participation in the system. The disadvantages, however, are also significant. People might transfer SCC to friends or family and build-up their CSCW without any real reason. It might also shatter the trust into a stable currency system. These considerations are the reason why a time-value of SCC is not part of the CSC system concept. To investigate the possibilities of the smart contracts on Ethereum, such a time-value was implemented. This mechanism leads to a deletion of all currency which is not spent after 30 days. This only affects SCC received as basic income and not as payment from others, which is achieved by individually registering every SCC coin and observing its history.

After the implementation and investigating various SCC exchanges on the system, we can draw several conclusions about the second research question and discuss what limitations and potentials blockchain technology offers for the CSC market system.

The main limitation we observed is the cost of running the system. Every function call of a smart contract must be compensated with ether (symbol ETH), which can be bought in exchange for EUR or USD. This money is used to compensate users who are providing computation power to the network, which encourages users to dedicate parts of their computation capacity in the first place. To prevent denial of service attacks, the cost of the function calls and transactions is proportional to the required computational effort.

The relation of cost to computational effort is an important aspect for the implementation of the social capital market system as an Ethereum smart contract. If the costs of a single transaction scale up proportionally to the number of users in the system, it would be too expensive to reasonably implement. This could be the case if the computations involve all previous transactions, e.g., for normalizing the CSCWs. Within the current system concept, this is not planned. Nevertheless, the cost, e.g., for the basic income is still significant.

This is closely related to another issue. As stated, all transactions require the payment of a fee paid in ETH. All market participants consequently need to buy ETH in order to participate in the market system. Financing a transaction with real money may limit the willingness of users to pay with SCC. On the other hand, this may also deter fraudulent behavior like pumping (see section 9.3).

A third limitation may be that transparency is a fundamental part of the blockchain. In the current design of the CSC system this is wanted and seen as an advantage. However, if a future extension of the system may want to implement, e.g., anonymous donations this would proof difficult on the Ethereum blockchain.

Another shortcoming that may become relevant for a large-scale role out of the system, is the limited smart contract execution speed. Currently, the Ethereum blockchain only allows to run up to five smart contracts per second, which limits the transaction speed if a huge number of people are using the system at the same time.

The fifth limitation is due to the nature of smart contracts: they are not supposed to be changed. As long as everything works well, this is an advantage. However, it is possible that unforeseen circumstances arise and make an intervention necessary, e.g., increasing or decreasing the size of α in equation 27 to adjust the speed of the CSCW build-up. This is not possible, and the only alternative is to switch to a new smart contract.

Besides the shortcomings, the experiment revealed also several chances of an implementation of the CSC market on the blockchain. The goals of the market system and the blockchain overlap in key areas, which leads to several advantages.

Firstly, the blockchain aims to create a safe environment for currency transactions, which is essential for the market system. It effectively prevents theft or double spending, i.e. people who try to spend the same currency twice. This level of security is difficult to implement in another way.

Secondly, the nature of the smart contracts complements the idea of a fair market. Smart contracts are the same for every participant, i.e. the same rules apply to everyone. This promotes transparency and fairness and thereby increases the trust in the system.

A third advantage is the created transparency. People are able to review previous transactions which makes the currency flows transparent and the overall system more trustworthy.

Part VI

SUMMARY OF RESULTS AND DISCUSSION OF IMPLICATIONS

The final part of this thesis consists of three chapters. Chapter 20 summarizes the results of the previous chapters and is structured along the overarching research questions. Based on these findings, critical conclusions about the potentials and limitations of a universal contributive social capital system are drawn in chapter 21. Finally, chapter 22 provides a brief outlook on potential future work that builds on or extends this research.

20

SUMMARY OF RESULTS

We investigated ways to assess a person's contributive social capital with the goal to improve online communication with the created transparency. Based on the design science approach by Hevner et al. (Hevner et al., 2004), we first motivated the problem with literature review and a social media survey with 242 participants (section 1.1). We identified several issues with online communication, e.g., the spread of fake news and lies and the difficulty to identify experts or to advertise own knowledge. Having transparency regarding the CSC of people in online interactions may help to mitigate these issues. Therefore, we envisioned a system that aims to create profound transparency with the help of several partly novel assessment approaches. Fundamental aspects regarding the main pillars of this system were investigated with a variety of experiments.

In this chapter we summarize the findings of these experiments and additional investigations along the initial research questions. Based on these findings, we can draw several conclusions about the potentials and limitations of contributive social capital systems. These conclusions are discussed in chapter 21 and point towards possible refinements of the concept.

- **Research question 1** (*What is a suitable personal characteristic that can be tracked and visualized to improve the experience of online communication?*)

  Section 2.1 suggested the personal characteristic "contributive social capital" as a metric to be assessed by the system. The CSC of a person encompasses their competence, trustworthiness, and social responsibility. By introducing a contextualization along different topics, as discussed in section 9.2.8, the transparency is increased further. This transparency regarding the CSC of people in online communication may act towards mitigation of the issues described in the beginning.

  We further discussed that CSC is related to other characteristics like reputation and influence and reviewed previous work on the assessment of these properties. The advantage of a single CSC value is that it can be assessed quickly and easily implemented in various applications (chapter 16). For the experimental investigations, the CSC ground truth was approximated by peer assessments that asked specifically about the three CSC constituents competence, trustworthiness, and social responsibility.

- **Research question 2** (*How could a system look like that automatically measures the contributive social capital of a person in a holistic way and maintains these scores over time?*)

  The concept of a system that measures the CSC of a person was presented in part II. There are several ways of how CSC can be measured (section 2.3). Thus, the system implements three different methods to create a holistic assessment. The first method is the analysis of previous interactions on online social networking and knowledge exchange platforms. This is complemented by a market system that automatically builds CSC weights based on the received currency. The currency transactions enable other goals, like the option to reward others for pro-social behavior or the introduction of new business models (section 5.1 and chapter 16). The system is completed by the ability to introduce CSC demonstrated offline into the system. This is achieved with certifications and endorsements.

There is a range of means for the assessment of individual characteristics, as discussed in chapter 10 and to some extend section 17.2. The advantages of the CSC system become apparent when compared to two other systems that aim to assess people based on their online behavior: Reddit's "karma system" and the Klout score. Similar to other threaded discussion and Q&A portals, the karma system of Reddit (section 3.5) builds a score for each user based on peer feedback about their individual contributions. The downside is that this score is only meaningful on the platform itself and that it is error prone and easily manipulated, e.g., by reposting popular contributions of others (Richterich, 2013). A system that is not limited to a single platform, is the Klout score (section 10.2.3). With supervised learning it creates an individual assessment of a user's influence in a way that is comparable to the first pillar of the CSC system. The goal and the usability of the Klout score, however, are different from the goal and applicability of the CSC system. Klout aims to identify influential people, which is useful for companies who want to hire influencers, e.g., for marketing purposes. The CSC system on the other hand aims to create transparency for everyone, not just paying customers and to promote pro-social behavior that benefits society overall.

None of the other systems implements a market system for the assessment. The use cases that require the ability to make payments (section 16) are consequently limited to the CSC system.

- **Research question 3** (*How can an individual CSC assessment for each user be extracted from previous interactions on OSNEM and other data sources?*)

  The extraction of CSC scores from different online sources is the first pillar of the CSC system. We investigated the assessment in five different experiments that are described in part III. The social capital experiment (chapter 11) provided a dedicated social networking platform to 242 participants and allowed them to interact in ways known from other OSNEM communities. With a CSC ground truth approximation, the extraction of CSC could be investigated with several supervised learning algorithms. Analogous investigations were performed in four additional experiments on excerpts from Facebook, Twitter, Quora, and a scientometrics data set (chapter 12).

  The experiments allow several conclusions about the assessment of CSC from different online data sources.

  In general, the findings support the hypothesis that contributive social capital can to some extent be extracted from previous interactions and contributions with the help of supervised learning algorithms. All algorithms were able to predict the individual CSC scores better than a random predictor. Rankings created with the predicted values and the ground truth approximations also correlated positively with the predictions in all experiments. This supports previous research about the extraction of other CSC related properties from online social networks (e.g., (Ziegler, 2009), (Yang et al., 2010), and (Su et al., 2012)).

The investigations were tailored to the respective platforms to reflect their individual properties. On Facebook we investigated 45 features for the analysis, 38 on Twitter, 44 in the scientometrics data set, and 41 on Quora. These features were selected based on previous work (chapter 10) and own investigations. Some of the features were created with NLP (e.g., topical similarities determined with LDA, or the SMOG measure of readability) or external input (e.g., university rankings).

The best algorithms for the prediction were random forests and neural networks. This indicates that the relationship between the ground truth assessment and the respective features cannot be described by a straightforward linear model.

The best results were achieved on data extracted from the Q&A platform Quora. This may be due to the content focus of Quora, where personal relationships are less important than, e.g., on Facebook. This indicates that similar platforms, like Stack Overflow or Reddit may also be ideal for CSC assessments.

The features that were of the highest importance for the prediction of CSCW were activity and feedback features. Activity features reflect the contributions of an individual to the overall network (e.g., number of posts on Facebook). Feedback features incorporate the responses by others (e.g., the number of "likes" a Facebook user receives). It is consequently important to use features from these categories for the CSC system. Features created from personal information of the users (e.g., the highest level of education of a Facebook user) or centrality measures calculated on interaction graphs (e.g., the betweenness centrality on the reply graph on Twitter) were generally of lower importance for the CSCW prediction.

The experiments also revealed limitations of the CSCW assessment. The accuracy of the prediction is, e.g, limited, which restricts the applicability of the CSC system as discussed in section 21.4. Additionally, the crawling process is time consuming. This may be countered by online assessments via the APIs of the platforms, once the algorithms are trained.

The potentials and limitations of a CSC assessment from online data sources are discussed in detail in section 21.1 in the following chapter.

- **Research question 4** (*How can a market system be used to build CSC scores of the market participants based on their transactions?*)

  The use of a market system for the assessment of personal characteristics is a novel approach that was investigated in part IV. The underlying idea is that a CSC weight is created based on received currency transactions in a dedicated market system. This hypothesis was tested with an experiment (chapter 14) and an agent-based simulation (chapter 15).

  The market experiment was part of the social capital experiment. All participants received an amount of virtual currency as basic income on a weekly basis.

Inspired by actions of other users, they were able to send and receive currency, e.g., to buy and sell information or to reward others for helpful contributions on the experiment's social networking platform.

Following the concept described in chapter 7, individual topic-specific CSCWs were created for each participant. A correlation analysis of the created weights with the respective ground truth assessments revealed a positive correlation between the corresponding values.

This supports the hypothesis that markets can be used for more than the mere transfer of currency, e.g., for the assessment of individual characteristics like CSC.

Section 21.2 in the following chapter summarizes the conclusions that can be drawn from the market experiments.

- **Research question 5** (*How is it possible to contextualize the CSC of a person along different topics?*)

  The contextualization of CSC by topics creates several new opportunities. It is the basis for expert identification as well as for advertising one's own knowledge and may help with the assessment of posts from otherwise unknown users.

  We investigated the question in two ways. At first we discussed how such a contextulization could look like and how it could be implemented. Secondly, we investigated the contextualization by categories experimentally during the market experiment.

  There are several options to divide CSC along categories that range from providing different "buckets" of equal standing to a structuring of topics along a hierarchy or graph (section 9.2.8). The latter two options have the advantage that CSC can propagate along them, i.e. increases in related topics can influence each other and thereby provide a more holistic view of an individual's interests. The respective topics can be organized in an ontology.

  In the market experiment, six topics were provided for the users. This led to the creation of six individual CSC scores and the identification of different experts in the topics (see chapter 14). The identified weights correlated with the respective ground truth CSC assessments, which indicates that a contextualization of CSCW along topics is possible.

- **Research question 6** (*How do participants experience the market system?*)

  To evaluate how the social capital market was perceived, we provided a survey to all participants. Overall, the responses were positive. 64% stated that they perceived SCC as more valuable than a "like" on Facebook because of its limited availability. The possibility to re-use it was appreciated by 61% who stated that receiving SCC was a positive feeling because of the possibility to reuse it. Over half of the participants (54%) would use SCC to pay for online news. The continuous voluntary participation over the whole duration of the experiment supports the hypothesis that the market was perceived well by the participants.

- **Research question 7** (*What are use cases of the CSC system for individuals, companies, and governmental institutions?*)

  Thanks to the monetary transactions on the market platform, the CSC system may create the basis for different use cases that go beyond the creation of transparency in online communication.

  First of all, it allows to assess people and their competence and trustworthiness in online interactions — as long as their CSCW is displayed on the respective platforms. Additionally, the CSCW can be used to identify experts in specific topics, which may be useful for individuals, companies, or even governmental institutions. With the help of subjectivity investigations, biases and impartial individuals can be identified. It also provides individuals with the opportunity to advertise their own skills and thereby effectively provide help to friends or family. Companies can also use the CSC system to improve their online businesses and use their CSCW as advertisement.

  The CSCW also lays the foundation for weight-based voting mechanisms. Weighting the vote of individuals by their CSCW may be used by groups to facilitate decision making or by companies to identify the best solution to problems over different departments.

  Finally, the system can be used to incentivize and support altruistic and prosocial behavior.

  These use cases are described in detail in chapter 16.

- **Research question 8** (*How can a CSC system be implemented in a secure and transparent way?*)

  Following the design science approach, we conceptualized the CSC system and implemented and refined its most important aspects. This includes the assessment of contributive social capital based on previous interactions and contributions to different online data sources and the build-up of CSC scores based on market interactions. The market implementation also included basic income and a selection of topics to distinguish between different categories.

  To implement the extraction mechanism on a large scale, user permissions are required to directly extract the required features from the platforms. The implementation of the market system may leverage blockchain technologies to create save and transparent transactions. The implementation as a smart contract on the Ethereum blockchain (chapter 19) revealed several issues, most importantly the compensation of the miners that leads to fees for all transactions. As this might deter people from using the system, a new CSC blockchain would be required for the implementation.

  All further design questions are described in section 9.

- **Research question 9** (*What are potentials and limitations of the CSC system?*)

  The main challenge in the context of the CSC system is to prevent misuse that is not in line with the goals of the system. Two examples of potential abuse were discussed in chapter 17: the harvesting of Facebook user data for political advertisements during the 2016 US presidential campaign, and the Chinese social credit system that penalizes unwanted behavior. While we suggest a variety of precautions to mitigate the risks of abuse, it is still important to keep these examples in mind to prevent them from happening.

Additional potentials and limitations of contributive social capital systems are discussed in detail in the following chapter 21.

- **Research question 10** (*Can the system be expanded to visualize subjectivity and biases of individuals regarding different topics?*)

  Visualizing the provenance of an individual's support, may allow an additional level of transparency regarding their motivation. We investigated ways to create this transparency in five experiments.

  For these investigations different data sets from Twitter, Facebook, a scientometrics database, Quora, and the social capital experiment were used. In all data sets we first identified different communities, i.e. groups of people who share commonalities. The people in the different clusters either shared topical interests, as identified with topic modeling, similar places of residency, or personal characteristics like gender.

  It was possible to visualize flows of CSC within and between these clusters, as well as on an individual level (chapter 18). We found that clusters with similar topical interests and backgrounds supported each other more, which may be an indication for homophily. The most inter-cluster support was observed on Quora, which is additional evidence that people on Quora are not limited to their own communities but interact along different topics of interest. On the individual level we visualized what percentage of a person's support comes from their own community and introduced ways to identify people who receive more or less support from a certain group than expected.

  These investigations can be used to identify users who are biased or impartial and directly transferred to different platforms and adjusted to cover the topics of interest.

The following chapter 21 draws conclusions about the potentials and limitations of CSC systems based on these findings and additional observations of the experiments.

21

CONCLUSIONS ABOUT THE POTENTIALS AND LIMITATIONS OF CSC SYSTEMS

The previous chapter 20 summarized the results of our experiments with regard to the initially stated research questions. Based on these findings and the discussions of the respective experiments, we can now draw several conclusions about the potentials and limitations of a contributive social capital system.

This chapter is structured along the four parts of the CSC system that were investigated in this thesis: extraction from OSNEM, modeling CSC scores with the market system, subjectivity investigations, and observations regarding the overall CSC system, which includes the system design.

21.1 POTENTIALS AND LIMITATIONS OF CSC EXTRACTION FROM ONLINE DATA SOURCES

The extraction of CSC scores from different online data sources was investigated in 5 experiments (chapter 11 and chapter 12). In these investigations we found several potentials of the CSC assessment approach:

- In all experiments it was possible to identify platform-specific features that could be used in the analysis. While all these features added some value to the prediction, some features are of lower importance. If required, these features may be omitted in the analysis to achieve two goals. Firstly, the time needed for data crawling and CSCW prediction is reduced and secondly, less features are needed which may be appealing to privacy sensitive users.

- In the social capital experiment (chapter 11) the CSCW predictions were about 20 percent better than a random predictor.

  The predicted values can also be used to create a ranking of the participants. This ranking correlated positively with a ranking based on the ground truth assessments. In the social capital experiment, the correlation coefficient was over 0.6 with a p-value below 0.001.

  These results indicate that it is generally possible to predict individual CSCWs for each participant based on the extracted features with the help of supervised learning. The prediction quality increased on the larger data sets extracted from Facebook, Twitter, Quora, and the citation network.

The experiments also indicated that the CSCW extraction from different data sources has some limitations:

- An important challenge is the amount of incorrect predictions. In all experiments there was at least one person whose CSCW was predicted incorrectly by the supervised learning algorithms. In the case of the social capital experiment, the predicted CSCW of 12 percent of the participants deviated by 25 or more percent from the ground truth assessment. That means 20 participants (12%) would have been represented wrongly within the system. Potential repercussions of these incorrect predictions are discussed in section 21.4. This issue may be addressed by combining the prediction with the assessments of the other pillars of the CSC system.

- Another challenge that all prediction tasks with supervised learning face, is determining a suitable ground truth value. Not all of our experiments could use an external ground truth assessment with the help of questionnaires. We investigated a variety of alternative ground truth approximations and discovered, based on previous work that there are some features that can be used as proxy. One observation was that the proxy should not be calculated based on too many other features, as they can no longer be used for the prediction task. Using a single feature that is carefully selected based on the results of the social capital experiment works well and can be used to compare different analysis methods or relations between features.

- It is also possible that the created CSC scores do not stop the spread of fake news or untrue information in all cases. The reason is that the CSCWs reflect many aspects of a person's character not only trustworthiness or truthfulness. This becomes clear at the hands of an example. For the CSCW prediction in the social capital experiment, the number of likes a person received was the most important feature, closely followed by the number and length of their comments. This indicates that a person who has many likes and writes many long comments has a high CSCW in the prediction. Received likes are certainly a measure for appreciation, and the number and length of comments are a measure for the activity of the user within the network. However, it is possible that some people react primarily to the entertainment value of a person's contributions rather than their truthfulness. This could mean that people who write many entertaining posts that do not necessarily reflect reality, receive many likes and consequently a large CSCW. To counter this issue, we used as many features as feasible for the analysis to create a holistic assessment. The combination of the predictions with input from the second and third pillar of the CSC system may additionally mitigate the risk of one-sided CSC assessments.

21.2 POTENTIALS AND LIMITATIONS OF CSC MODELING WITH CSC MARKETS

Modeling CSC scores with the help of market interactions was investigated in two experiments. The first experiment, described in chapter 14, provided all participants with a platform on which they could trade virtual currency. The second experiment was a market simulation with 1,000 agents. Based on these investigations we can draw several conclusions about the potentials and limitations of market-based assessments.

The experiments revealed several promising results:

- The predicted CSCWs in the experiment and the simulation correlated positively with the ground truth assessments.

- A contextualization by topics was possible and the topic-specific CSC assessments also correlated positively with the respective ground truth assessments. The correlation coefficients of the different predictions were between 0.21 and 0.41 which indicates a moderate correlation with p-values below 0.001.
- The market was perceived as useful by the participants and especially the usability of the currency was appreciated by most of the users as assessed with a survey of 115 participants (see section 14.5.1).
- A variety of use cases for the market system could be identified. They range from using the currency to pay for online news to offering services with SCC as payment. Receiving SCC increases the CSCW of the person which may, e.g., be used to advertise their skills (see section 16).
- The market offers an easy way to reward pro-social behavior in the form of SCC transactions. 71% of 242 students said in our social media survey (section 16.7) that they are willing to use this mechanism to reward others. Previous research suggests that such rewards will further promote social behavior (Gneezy et al., 2011) and thereby achieve one of the goals of the CSC system.

The experiments also revealed several limitations of the market assessment:

- The main challenge of the market system is similar to the previous section: the prediction has an error and there is and likely always will be at least one person whose score is predicted incorrectly. In the case of the social capital market experiment most people did not receive SCC transactions in all topics, i.e. they have at least one weight that is still at the starting value of 1. This reflects that the person did not interact in this topic during the experiment. Unless the person really has no CSC in the topic, this also means that most people are represented wrongly in at least one of the six categories. Depending on the use case, this may have serious consequences, as discussed in section 21.4.
- In both experiments the average transaction price increased over time. This may be interpreted as inflation due to the increase of the overall available currency. If the CSC system is rolled out without adjustments, it is likely that it also includes inflation. As a result, prices would need to be adjusted on a regular basis and people might lose faith in the system. We described ways to limit the experienced inflation, e.g., by introducing trade taxes.
- The CSCWs in the market system are built and maintained based on received transactions of virtual currency. The underlying hypothesis is that people will receive currency due to help and services they provided to others. This has been backed up to some extent by the correlation analysis of the market system experiment. However, it is possible that entrepreneurial endeavors of individuals lead to an accumulation of SCC and thereby CSCWs based on characteristics like entrepreneurship. It needs to be investigated in future work if this would impact the system in a negative way.

- Another shortcoming is inequality regarding the distribution of SCC and CSCW. This is partly wanted because the goal of the system is to reward helpful contributions and people who add value to the network. However, the perceived inequality may displease participants, which needs to be investigated from a psychological point of view. If the perceived inequality discourages people, it may be necessary to investigate a refinement of the market distribution parameters as described by equations 20, 21, 22, and 23.

21.3 POTENTIALS AND LIMITATIONS OF SUBJECTIVITY INVESTIGATIONS WITH THE CSC SYSTEM

Some people on online social networking platforms and even in academic citation networks have a high CSCW solely because of support from one group of people. While this high CSCW may still be well-earned, having transparency about the origin of the CSCW of individuals or groups, may provide additional insights about their motivation. We conducted five experiments on Twitter, Facebook, Quora, a scientometrics data set, and the social capital experiment data to investigate measures to assess subjectivity and biases and ways to visualize them.

The potentials identified in these investigations are:

- It is possible to identify communities of people who represent different sides of debates. This task was achieved with graph clustering or by using individual characteristics or joint interests to group the users.
- Once clusters are identified it is possible to visualize the flows of CSC within and between the different clusters. The subjective support matrix that was introduced in section 18.3 can be used for the visualization.
- The visualization of subjectivity also works on an individual level: People who belong to one of two groups can be flagged if they receive more or less support from one of the groups than expected. The reason for these deviations need to be examined individually and allow to identify biased or impartial individuals.

The experiments also revealed some limitations of the approach:

- The clustering of the users proved difficult in some instances. While attribute-based clustering, e.g., according to country of origin is straightforward, it may be difficult to identify meaningful clusters with graph-clustering or topic based clustering. The reason is that these clusters often display an overlap and need to be interpreted individually. In many cases this is no problem and may even provide additional insights, e.g., about whether there are well-defined interest groups. It may, however, complicate the subjectivity analysis along a specific topic of interest. In this case additional classifiers may be required, e.g., to predict the political orientation of users on Facebook (David et al., 2016).
- The subjectivity analysis can only be conducted if appropriate data is available. On Facebook, e.g., it requires all of the user's previously received likes as well as the whole surrounding network including their post history to assess thematic interests.

- So far we have only discussed the created transparency in a positive light. However, it may be possible that the knowledge from where a person receives support is misused. Seeing that an individual is mainly supported from a group whose values another person does not share may, e.g., lead to misunderstandings.

- Finally, the precision of the analysis does not allow to display subjectivity correctly in all cases. While the visualization of CSC flows is precise, the interpretation of the different clusters may not always be possible. Individuals may also be assigned to a cluster that they may not identify with.

21.4 POTENTIALS AND LIMITATIONS OF THE OVERALL CSC SYSTEM

This section summarizes the chances and limitations of the CSC system as a whole with regard to the system architecture and additional investigations.

Our investigations revealed the following potentials:

- Literature reviews and a social media survey that we conducted indicates that there is a need for more transparency in online communication. This is indicated by several issues: the spread of fake news (Goodwin-Ortiz, 2017) and falsehoods (Wang and Zhuang, 2018), wrong self-assessments (Anson, 2018) and deliberate deceptive behavior on social media (Underwood et al., 2011), which may lead to depression (Sidani et al., 2016). Additionally, there were reports of data misuse from the side of the platform providers (NY-Times, 2018). This was backed up by the study participants who found it difficult to identify fake news, find experts, or advertise their own expertise. The CSC system described in part II was designed to address these issues.

- The key parts of the CSC system are the analysis of online data sources and the market assessment. As described in the previous sections, both pillars could be used for the assessment of CSCWs of individual participants in several experiments. Additionally, essential parts of the system were perceived positively by the participants (see section 14.5.1) and we identified several use cases and business models that can be created with the system (chapter 16).

- It is possible to implement the market system as a smart contract on the Ethereum blockchain. This can be seen as a proof of concept regarding the implementation of a market system in a secure and transparent way.

However, the experiments also revealed several limitations:

- A major challenge of the system was already mentioned in the previous sections: the limited prediction accuracy. While it may still be favorable to have a system that occasionally predicts a wrong expert when looking for an answer, it may be much more serious when the system is used for other applications. This becomes apparent when looking at two different systems that are already in operation.

  The Klout score computes a measure that describes a person's media influence, as discussed in section 10.2.3. In the past, this score has been used by recruiters in the HR departments of companies when hiring people for marketing positions (Forbes, 2012). If the CSC system were to be used for similar purposes – which is not in the hands of the researchers implementing the system – a wrong assessment may consequently cost someone their job.

Even more serious are the consequences of incorrect predictions in the case of the Chinese social credit system (section 17.2). It implements punishments like travel bans for people who receive low scores (TheTelegraph, 2017). While the CSC system design implements methods to prevent the CSC system from being used in this way, it underlines the severity of potentially wrong assessments.

A combination of all three assessment pillars may improve the limited prediction quality of the individual pillars.

- The scale on which the CSCW is displayed also revealed limitations of the system. In section 9.2.9, several scaling options were discussed. In the social capital and market experiment (chapter 11 and chapter 14) a linear scaling was sufficient. In the market simulation (chapter 15), a linear scaling would have led to CSCWs of up to 100,000. As demonstrated, logarithmic scaling prevents such unpractical values.

To take full advantage of the chances of the CSC system, the limitations discussed in this and the previous sections need to be addressed in the suggested ways. Based on the experiences in this design science circle, we suggest several new experiments that may suitable next steps. These experiments are listed in the following, final chapter 22.

22

FUTURE WORK

As described in section 1.3 this thesis partly follows the design science approach. This work can be seen as part of a first iteration of a design science circle. We created a concept of a system with the goal to tackle several issues that are present in modern online communication. This concept was tested with the help of several experiments, as described in chapters 11, 12, 14, 15, and 18. The analyses and evaluations of the experiments have created further insights on the workings of the system. These insights were summarized with regard to the initial research questions in chapter 20. Based on these findings, we could draw several conclusions about the potentials and limitations of a contributive social capital system (chapter 21) that point towards possible refinements of the system. With the help of additional experiments a new design science cycle can be initiated. Ideas for future work are discussed in this chapter to facilitate the design of new experiments.

- **Large scale experiment.** We conducted an experiment with 242 participants to evaluate the first and second pillar of the CSC system. This experiment provided valuable insights and acts as a proof of concept. To validate these findings on a larger scale, excerpts from other social networking platforms and a scientometrics database were used. These data sets entail contributions created over several years and were extracted from Facebook (11,629 users), Twitter (25,000 users), Quora (3,069 users), and the scientific platform ArnetMiner (99,178 authors). While these data sets were much larger, they lacked a precise ground truth CSC assessment.

  To counter these shortcomings, we suggest a large scale, long term experiment with at least 1,000 users who interact on their preferred social media platforms and a market system and who provide additional ground truth assessments. This may enable further fine-tuning of the market parameters with regard to the perceptions by different user groups.

- **Combination of the three pillars.** In the course of our research, no experiment about certifications and endorsements was conducted. The main reason is that these assessments are more exciting from a psychological and sociological than a computer science perspective. Once such additional experiments are conducted, it is important to investigate how the three pillars can be combined. In section 9.1 the basis for the combination is discussed. The scaling of each of the three assessments cannot be determined by theory only, but requires field studies to ensure that the resulting score is indeed a correct representation of the person's CSC.

- **Ontology design.** The market experiment allowed the users to select one of six topics for their SCC transactions. While this was successful, a more comprehensive topic selection is required for future large-scale implementations of the CSC system. As described in section 9.2.8.2 there are several ways to achieve this, including an ontology that connects related topics in graphs or hierarchical structures (e.g., "astrophysics" belongs to "physics" and "physics" belongs to "science"). It is up to future experiments to evaluate the different options and define how CSC received in one topic propagates to related categories.

- **Deep dive market experiment.** The market system offers additional areas for investigations besides the large-scale implementation. The focus of our experiments was to examine its basic functionality. New experiments can build on this mechanism and introduce additional properties like trade taxes. We hypothesized that introducing taxes may limit the observed inflation and be used as a means to finance the basic income. This needs to be investigated with a focus on the psychological implications for the users of the system. Different means of fraud and their countermeasures, as described from a theoretical point of view in section 9.3, are also interesting fields for future research.

Once these investigations enable a fine-tuning of the system parameters with regard to the users' perception of the system, the CSC system can be implemented on a full scale.

We hope that this work provides a helpful basis for future investigations and that all other scientists working on this topic are as thrilled by the idea of improving online communication for everyone as we are.

Part VII

APPENDIX

Appendix A lists all questions that were used to assess the CSC ground truth value of the participants of the social capital experiment (chapters 11 and 14). A selection of post titles is provided in appendix B, which exemplifies what topics were discussed in the experiment. The interface of the market experiment is depicted and explained in appendix C. Appendix D lists all bachelor's and master's theses and guided research projects that were supervised in the course of this thesis. Finally, appendix E shows a complete list of all excerpts from the own publications that were directly used in this thesis.

A

GROUND TRUTH SURVEY

In the experiment, the participants could assess the competence, trustworthiness, and social responsibility of other known users with the help of eight questions. All assessments were made on a scale from 0 to 100:

A.1 COMPETENCE ASSESSMENT

Please evaluate this person's competence (mixture of knowledge and expertise) in the following three fields:

- Populism in politics (e.g., Trump's wall to Mexico, refugee crisis in Europe, etc.)
- Living in Munich (e.g., sports and leisure activities, finding affordable living, lectures at TUM, etc.)
- Healthy food and sustainability (e.g., calorie counts, genetically altered nutrition, sustainability, etc.)

A.2 TRUST ASSESSMENT

Please evaluate how much you trust this student.

- What is your general level of trust towards this student?
- To what extent is this person concerned for your welfare — someone who is looking out for you, who would go out of their way to help you, and who would not knowingly do anything to hurt you?
- To what extent is this person fair and honest — do they stick to their word and use sound principles to guide themselves?

A.3 SOCIAL RESPONSIBILITY

Please help us understand how environmentally friendly and socially engaged the selected person is.

- Environmental friendliness (e.g., support of environmental protection institutions, sustainable food, waste separation, etc.)
- Social support/engagement (e.g., support of friendly societies, help to other students/friends/strangers, support for elderly family members, etc.)

B

SELECTION OF POST TITLES ON THE SOCIAL NETWORKING PLATFORM

To facilitate a better understanding of the happenings in the social networking platform, this chapter provides a selection of the discussed topics. The headlines of selected blog posts are listed without censoring or correcting grammar. The posts were written by students and are often related to events that took place in summer 2017, e.g., Trump's presidency or places to escape the summer heat.

B.1 EXEMPLARY 'POPULISM IN POLITICS' DISCUSSION TOPICS

- Twittering presidents
- Politics in social networks
- Macron - France' new monarch?
- Trump exits Paris' climate agreement
- Gay marriage to be legalized in Germany following today's vote
- The fuck is going on in Hamburg
- Trump and his attitude towards women
- Political quiz
- Do you think the current events regarding the President will result in new elections?
- Death penalty, aka capital punishment in the US
- Trump's final mistake
- Trump: I shared information with Russia and I had absolute right to do so - Guardian
- Is Obama a reptile?
- Is Obama amphibian?

B.2 EXEMPLARY 'LIVING IN MUNICH' DISCUSSION TOPICS

- The 13 beer gardens everyone in Munich should know
- Accommodation in Munich
- Hiking around Munich

- Craft beer bar-shop in Munich
- Weekend getaways near Munich
- What is your favorite lake near Munich?
- Sakura cherry blossom Munich
- Go kart in Munich
- What is the best secret attraction in Munich?
- Comedy clubs in Munich
- Best places in Munich to enjoy the sun and have a swimbath
- Best cinema in Munich
- Sightseeing in Munich
- Best places for brunch in Munich
- Road cycling in Munich
- Exploring Munich
- Living in Munich - How much money do students need per month?
- Living in Munich - What to show visitors?
- Places to escape the heat
- Best lasertag in Munich

B.3 EXEMPLARY 'HEALTHY FOOD AND SUSTAINABILITY' DISCUSSION TOPICS

- Meat lover's guide to healthy eating
- Living sustainable - is it too expensive?
- Spicy food
- Healthiness of food
- Meat is murder - why nobody cares
- Who knows the best burrito in Munich?
- Cooking food - where do you find your recipes?
- Living healthy - is it too expensive?
- Eating health and weight loss
- Pizza
- Does anyone grow vegetables in their flat?

B.4 EXEMPLARY DISCUSSION TOPICS ABOUT THE LECTURE AND THE EXPERIMENT

- Script for downloading all resources from the Piazza class page
- We need a downvote button
- No emojis allowed
- iPyhton notebooks and 40 minutes
- Talking about data analysis
- Exams and grades
- What working with pandas feels like
- Sort function for the market place
- Credit transfer
- How do I create a blog?
- The coolest Prof I know is Georg Groh
- Everybody can share their opinion - unfortunately everybody does
- Board rules

C USER PROFILE IN MARKET SYSTEM

The social capital market allows users to send and receive currency (SCC). In the top row of the interface seven buttons are displayed. The **Home** button leads to the starting page, where a basic introduction to the functionality of the system as well as news are shown. **Settings** allows the participant to review their information. The **Help & Contact** section provides the email addresses of Sebastian Schams and Georg Groh, as well as the contact of the teaching assistant who supervised the experiment.

The **Credit Transfer** side is displayed in the screenshot in figure 27. It shows the current account balance (here: 200 Social Capital Currency) and a list of users with their profile pictures.

Clicking on a user opens a new window in which a transaction can be specified (including the amount, a topic, and a personal message). The **User Survey** asked a variety of questions regarding the whole experiment and the experience of the users.

In the section **Questionnaire**, users could provide feedback about the ground truth contributive social capital of others (see appendix A).

All inputs were saved in a SQLite database.

Figure 27: Screenshot of the market exchange platform used for the social capital market experiment described in chapter 14.

D

SUPERVISED THESES AND RESEARCH PROJECTS

During the course of this thesis a variety of bachelor's and master's theses, as well as one guided research project were supervised. A complete list is provided in table 25.

Student	**Title of project or thesis**
Simon Zettler	Evaluation of Concepts for Social Capital Integration in Social Networks (Zettler, 2016)
David Hauer	Design of a Smart Contract Based Social Capital Currency System on Ethereum (Hauer, 2016)
Monika Varshney	Models of Topic-Specic Reputation / Social Capital in Online Social Networks (Varshney, 2017)
Christian Höfer	Designing an experiment for the evaluation of social capital assessment and market concepts (Höfer, 2017)
Valeriia Chernenko	Subjectivity and Social Capital Scores Determination in Threaded Communication Platforms (Chernenko, 2017)
Johannes Feil	Subjectivity and Social Capital Scores Determination in Microblogging Platforms and Citation Networks (Feil, 2017)
Panagiota Revithi	Design of Social Capital Markets (Revithi, 2017)
Maximilian Schmidt	Social Capital Market Simulation (Schmidt, 2017)
Rauf Zeynalov	Cross-medial Machine Learning based Social Capital Analysis (Zeynalov, 2018)
Patrick Süß	Evaluation of Social Capital Markets and Discussion of Use-Cases (Süß, 2018)
Niclas Hirtle	Social Capital Market Simulation in Consideration of Macroeconomic Concepts (Hirtle, 2018)
Daniel Clot	Investigating Ontology Design for Contributive Social Capital Assessments (Clot, 2018)

Table 25: Supervised research projects, bachelor's, and master's theses

E

LIST OF DIRECT EXCERPTS FROM OWN PUBLICATIONS

Several articles were published in the context of this thesis. This thesis contains direct excerpts from these papers. A complete list of all excerpts is provided in the following. In some cases, the excerpts were slightly altered (e.g., the references were adjusted to fit the thesis).

CHAPTER 2

Section 2.1.1 (first seven sentences) is from (Schams et al., 2018a) and section 2.2 (all subsections) is from (Schams and Groh, 2018)

CHAPTER 3

Section 3.1 (first paragraph and the first sentence of the second paragraph), section 3.2 (first paragraph), section 3.3 (first paragraph), section 3.4 (first, second, and third paragraph), and section 3.5 (first paragraph and the first sentence of the second paragraph) are from (Schams and Groh, 2018)

CHAPTER 5

Section 5.2 (first paragraph), section 5.3 (second sentence), and figure 6 are from (Schams et al., 2018a)

CHAPTER 6

Section 6.2 is from (Schams et al., 2018a)

CHAPTER 7

Section 7.1 (second and third sentence), section 7.2 (all but the last three sentences), section 7.3 (first paragraph and second and fourth sentence of subsection 7.3.3) are from (Schams et al., 2018a)

CHAPTER 10

All subsections with the exception of the sentences that reference the CSC system are from (Schams and Groh, 2018)

CHAPTER 11

Section 11.3 (eighth paragraph without the first two sentences and ninth paragraph without the first sentence), section 11.4 (second and third sentence, second paragraph, subsection 11.4.1 without the last sentence, subsection 11.4.2, subsection 11.4.3, subsection 11.4.3), section 11.5 (first paragraph, subsection 11.5.2, and subsection 11.5.3), section 11.6, section 11.7, and section 11.8 (third to sixth sentence) are from (Schams et al., 2018a)

CHAPTER 13

Section 13.1 (second and third paragraph) is from (Schams et al., 2018b)

CHAPTER 14

Section 14.3 (itemization and the following sentence), section 14.4 (second and third paragraph), section 14.5 (first sentence, subsection 14.5.1 except the last sentence, subsection 14.5.2, and subsection 14.4.3), figure 15, table 15, table 16, and table 17 are from (Schams et al., 2018b)

CHAPTER 15

Section 15.3, section 15.5 (all but the last paragraph), section 15.6 (itemization), section 15.7 (last sentence), figure 16, figure 17, figure 18, and figure 19 are form (Schams et al., 2018b)

APPENDIX A

All sections are from (Schams et al., 2018a).

BIBLIOGRAPHY

Abbas, A. and Khan, S. U. (2014). A review on the state-of-the-art privacy-preserving approaches in the e-health clouds. *IEEE Journal of Biomedical and Health Informatics*, 18(4):1431–1441.

Abbasi, A., Wigand, R. T., and Hossain, L. (2014). Measuring social capital through network analysis and its influence on individual performance. *Library & Information Science Research*, 36(1):66–73.

Aggarwal, C. (2011). *Social Network Data Analytics*. Springer US.

Aker, A., Kurtic, E., Balamurali, A., Paramita, M., Barker, E., Hepple, M., and Gaizauskas, R. (2016). A graph-based approach to topic clustering for online comments to news. In *European Conference on Information Retrieval*, volume 9626, pages 15–29. Springer.

Alpaydin, E. (2010). *Introduction to Machine Learning*. Adaptive computation and machine learning. MIT Press.

Alpaydin, E. (2016). *Machine Learning: The New AI*. The MIT Press Essential Knowledge series. MIT Press.

Anger, A. and Köhler, J. (2010). Including aviation emissions in the eu ets: Much ado about nothing? a review. *Transport Policy*, 17(1):38–46.

Anger, I. and Kittl, C. (2011). Measuring influence on Twitter. *Proceedings of the 11th International Conference on Knowledge Management and Knowledge Technologies - i-KNOW '11*, page 1.

Anson, I. (2018). Partisanship, Political Knowledge, and the Dunning-Kruger Effect. *Political Psychology*.

Aspers, P. (2011). *Markets*. Economy and Society. Wiley.

Awad, M. and Khanna, R. (2015). *Efficient Learning Machines: Theories, Concepts, and Applications for Engineers and System Designers*. Apress.

Balog, K. (2006). Finding experts and their Details in e-mail corpora. *of the 15th international conference on World*, (i):3–4.

Bentwood, J. (2008). Distributed Influence : Quantifying the Impact of Social Media. *http://www.crmxchange.com/uploadedFiles/White_Papers/PDF/edelman-white-paper-distributed-influence-quantifying-the-impact-of-social-media.pdf*, pages 1–16. (accessed: 20-01-2018).

Bergstra, J., Yamins, D., and Cox, D. D. (2013). Hyperopt: A python library for optimizing the hyperparameters of machine learning algorithms. In *Proceedings of the 12th Python in Science Conference*, pages 13–20. Citeseer.

Bergstrom, K. (2011). "don't feed the troll": Shutting down debate about community expectations on reddit.com. *First Monday*, 16(8).

Bertot, J., Jaeger, P., and Grimes, J. (2010). Using icts to create a culture of transparency: E-government and social media as openness and anti-corruption tools for societies. *Government information quarterly*, 27(3):264–271.

Best, S. and Krueger, B. (2006). Online interactions and social capital: Distinguishing between new and existing ties. *Social science computer review*, 24(4):395–410.

Bhattacharjee, R. and Goel, A. (2005). Avoiding Ballot Stuffing in eBay-like Reputation Systems. *Proceedings of the 2005 ACM SIGCOMM Workshop on Economics of Peer-to-peer Systems*, pages 133–137.

Bird, C., Gourley, A., Devanbu, P., and Gertz, M. (2006). Mining Email Social Networks. *MSR '06 Proceedings of the 2006 international workshop on Mining software repositories*, pages 137–143.

Biswas, D., Biswas, A., and Das, N. (2006). The differential effects of celebrity and expert endorsements on consumer risk perceptions. the role of consumer knowledge, perceived congruency, and product technology orientation. *Journal of Advertising*, 35(2):17–31.

Blei, D., Carin, L., and Dunson, D. (2010). Probabilistic topic models. *IEEE signal processing magazine*, 27(6):55–65.

Blei, D., Ng, A., and Jordan, M. (2003). Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan):993–1022.

Bohn, A., Buchta, C., Hornik, K., and Mair, P. (2014). Making friends and communicating on facebook: Implications for the access to social capital. *Social Networks*, 37:29–41.

Boshmaf, Y., Beznosov, K., and Ripeanu, M. (2013). Graph-based sybil detection in social and information systems. *Proceedings of the 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, pages 466–473.

Bouguessa, M. and Ben Romdhane, L. (2015). Identifying Authorities in Online Communities. *Acm Transactions on Intelligent Systems and Technology*, 6(3):23.

Bozzon, A., Brambilla, M., Ceri, S., Silvestri, M., and Vesci, G. (2013). Choosing the Right Crowd: Expert Finding in Social Networks. *Proceedings of the 16th International Conference on Extending Database Technology - EDBT '13*, pages 637–648.

Breiman, L. (2001). Random forests. *Machine Learning.*, 45(1):5–32.

Breiman, L. (2017). *Classification and Regression Trees*. CRC Press.

Brown, J., Broderick, A., and Lee, N. (2007). Word of mouth communication within online communities: Conceptualizing the online social network. *Journal of interactive marketing*, 21(3):2–20.

Bry, F., Ebner, M., Pohl, A., and Taraghi, B. (2014). Interaction in Massive Courses J. UCS Special Issue. *Journal of Universal Computer Science*, 20(1):1–5.

Burke, M., Kraut, R., and Marlow, C. (2011). Social capital on facebook: Differentiating uses and users. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '11, pages 571–580, New York, NY, USA. ACM.

Buterin, V. et al. (2014). A next-generation smart contract and decentralized application platform. *White paper*.

Campbell, C., Maglio, P., Cozzi, a., and Dom, B. (2003). Expertise identification using email communications. *Cikm 2003*, (January):528–531.

Carmel, D., Roitman, H., and Yom-Tov, E. (2012). On the relationship between novelty and popularity of user-generated content. *ACM Trans. Intell. Syst. Technol.*, 3(4):69:1–69:19.

Cha, M., Haddadi, H., Benevenuto, F., and Gummadi, K. (2010). Measuring user influence in twitter: The million follower fallacy. *4th International AAAI Conference on Weblogs and Social Media (ICWSM)*, pages 10–17.

Cheng, C. and Shuyang, O. (2014). The status quo and problems of the building of china's social credit system and suggestions. *International Business and Management*, 8(2):169–173.

Cheng, R. and Vassileva, J. (2006). Design and evaluation of an adaptive incentive mechanism for sustained educational online communities. *User Modeling and User-Adapted Interaction*, 16(3):321–348.

Chernenko, V. (2017). Subjectivity and Social Capital Scores Determination in Threaded Communication Platforms. Master's thesis, *Technical University of Munich*, Germany. Supervised by Sebastian Schams and Georg Groh.

Christofides, E., Muise, A., and Desmarais, S. (2009). Information disclosure and control on Facebook: Are they two sides of the same coin or two different processes? *Cyberpsychology & behavior*, 12(3):341–345.

CIA (2016). Distribution of family income - Gini index. *https://www.cia.gov/library/publications/the-world-factbook/fields/2172.html* (accessed February 22, 2018).

Clauset, A., Newman, M., and Moore, C. (2004). Finding community structure in very large networks. *Physical review E*, 70(6):066111.

Clot, D. (2018). Investigating Ontology Design for Contributive Social Capital Assessments. Master's thesis, *Technical University of Munich*, Germany. Supervised by Sebastian Schams and Georg Groh.

CNBC (2017). Cash is already pretty much dead in China as the country lives the future with mobile pay. *https://www.cnbc.com/2017/10/08/china-is-living-the-future-of-mobile-pay-right-now.html* (accessed June 21, 2018).

Corcho, O., Fernández-López, M., and Gómez-Pérez, A. (2003). Methodologies, tools and languages for building ontologies. where is their meeting point? *Data & knowledge engineering*, 46(1):41–64.

da Silva, I., Spatti, D., Flauzino, R., Liboni, L., and dos Reis Alves, S. (2016). *Artificial Neural Networks: A Practical Course*. Springer International Publishing.

Danezis, G. and Mittal, P. (2009). Sybilinfer: Detecting sybil nodes using social networks. In *NDSS*, pages 1–15. San Diego, CA.

Dannen, C. (2017). *Introducing Ethereum and Solidity*. Springer.

David, E., Zhitomirsky-Geffet, M., Koppel, M., and Uzan, H. (2016). Utilizing Facebook pages of the political parties to automatically predict the political orientation of Facebook users. *Online Information Review*, 40(5):610–623.

Davidson, P. (2003). *Financial Markets, Money, and the Real World*. Edward Elgar Publishing, Incorporated.

Dellarocas, C. (2003). The digitization of word of mouth: Promise and challenges of online feedback mechanisms. *Management Science*, 49(10):1407–1424.

Dinger, J. and Hartenstein, H. (2006). Defending the sybil attack in P2P networks: Taxonomy, challenges, and a proposal for self-registration. In *The First International Conference on Availability, Reliability and Security, 2006. ARES 2006.*, pages 756–763. IEEE.

Douceur, J. R. (2002). The sybil attack. In *International workshop on peer-to-peer systems*, pages 251–260. Springer.

Earnest & Young (2017). Online-Nutzung in Deutschland. *http://www.ey.com/Publication/vwLUAssets/ey-online-nutzung-in-deutschland-juni-2017/$FILE/ey-online-nutzung-in-deutschland-juni-2017.pdf* (accessed June 1, 2018).

Egghe, L. (2006). Theory and practise of the g-index. *Scientometrics*, 69(1):131–152.

Ellerman, A. D. and Buchner, B. K. (2007). The european union emissions trading scheme: Origins, allocation, and early results. *Review of Environmental Economics and Policy*, 1(1):66–87.

EuropeanCommision (2010). The EU ETS is delivering emission cuts. *https://ec.europa.eu/clima/sites/clima/files/docs/factsheet_ets_emissions_en.pdf* (accessed June 22, 2018).

Everling, O. (2013). *Credit Rating durch internationale Agenturen: eine Untersuchung zu den Komponenten und instrumentalen Funktionen des Rating*. Springer-Verlag.

Faliagka, E., Tsakalidis, A., and Tzimas, G. (2012). An integrated e-recruitment system for automated personality mining and applicant ranking. *Internet research*, 22(5):551–568.

Feil, J. (2017). Subjectivity and Social Capital Scores Determination in Microblogging Platforms and Citation Networks. Master's thesis, *Technical University of Munich*, Germany. Supervised by Sebastian Schams and Georg Groh.

Fenton, N. and Neil, M. (2012). *Risk Assessment and Decision Analysis with Bayesian Networks*. CRC Press.

Fernandez-Cano, A., Torralbo, M., and Vallejo, M. (2004). Reconsidering price's model of scientific growth: an overview. *Scientometrics*, 61(3):301–321.

Forbes (2012). Will Your Klout Score Get You Hired? The Role of Social Media in Recruiting. *https://www.forbes.com/sites/jeannemeister/2012/05/07/will-your-klout-score-get-you-hired-the-role-of-social-media-in-recruiting/#8020e58e64b0* (accessed May 25, 2018).

Fukuyama, F. (2001). Social capital, civil society and development. *Third World Quarterly*, 22:7–20.

Gabbriellini, S. (2014). The evolution of online forums as communication networks: An agent-based model. *Revue française de sociologie*, 55(4):805–826.

Gambetta, D. (1988). Can we trust trust? In *Trust: Making and Breaking Cooperative Relations*, pages 213–237. Basil Blackwell.

Garcia, S. M., Tor, A., and Schiff, T. M. (2013). The psychology of competition: A social comparison perspective. *Perspectives on Psychological Science*, 8(6):634–650.

Garrett, R. K. (2009). Echo chambers online?: Politically motivated selective exposure among internet news users. *Journal of Computer-Mediated Communication*, 14(2):265–285.

Gelman, A. and Hill, J. (2006). *Data Analysis Using Regression and Multilevel/Hierarchical Models*. Analytical Methods for Social Research. Cambridge University Press.

Gilbert, E. (2013). Widespread Underprovision on Reddit. *Proceedings of the 2013 conference on Computer-supported Cooperative Work*, pages 803–808.

Gini, C. (1912). Variabilità e mutabilità. *Reprinted in Memorie di metodologica statistica (Ed. Pizetti E, Salvemini, T). Rome: Libreria Eredi Virgilio Veschi*.

Gjoka, M., Kurant, M., Butts, C. T., and Markopoulou, A. (2009). A walk in facebook: Uniform sampling of users in online social networks. *CoRR*, abs/0906.0060.

Gneezy, U., Meier, S., and Rey-Biel, P. (2011). When and why incentives (don't) work to modify behavior. *Journal of Economic Perspectives*, 25(4):191–210.

Golbeck, J. (2008). Weaving a web of trust. *Science*, 321(5896):1640–1641.

Golbeck, J. (2009). *Computing with Social Trust*. Springer.

Golbeck, J., Robles, C., and Turner, K. (2011). Predicting personality with social media. In *CHI'11 extended abstracts on human factors in computing systems*, pages 253–262. ACM.

Gómez, V., Gómez, V., Kaltenbrunner, A., Kaltenbrunner, A., López, V., and López, V. (2008). Statistical analysis of the social network and discussion threads in slashdot. *Proceeding of the 17th international conference on World Wide Web - WWW '08*, page 645.

Gómez-Pérez, A. and Benjamins, R. (1999). Overview of knowledge sharing and reuse components: Ontologies and problem-solving methods. IJCAI and the Scandinavian AI Societies. CEUR Workshop Proceedings.

Gomez-Uribe, C. A. and Hunt, N. (2016). The netflix recommender system: Algorithms, business value, and innovation. *ACM Transactions on Management Information Systems (TMIS)*, 6(4):13.

Goodwin-Ortiz, C. (2017). Fake News On Social Media: Illusory Truth and the 2016 Presidential Election. Bachelor's thesis, *The State University of New Jersey*, USA.

Granovetter, M. (1976). Network sampling: Some first steps. *American journal of sociology*, 81(6):1287–1303.

Groh, G. and Birnkammerer, S. (2011). Privacy and information markets: Controlling information flows in decentralized social networking. In *Privacy, Security, Risk and Trust (PASSAT) and 2011 IEEE Third Inernational Conference on Social Computing (SocialCom), 2011 IEEE Third International Conference on*, pages 856–861. IEEE.

Guarino, N., Oberle, D., and Staab, S. (2009). What is an ontology? In *Handbook on ontologies*, pages 1–17. Springer.

Hadgu, A. T. and Jäschke, R. (2014). Identifying and analyzing researchers on twitter. In *Proceedings of the 2014 ACM conference on Web science*, pages 23–32. ACM.

Hartl, F. (2013). Topic Recommender Systems in Social Networks Using Topic Models. Master's thesis, *Technical University of Munich*, Germany. Supervised by Jan Hauffa and Georg Groh.

Hassan, S. (2013). Identifying criteria for measuring influence of social media. *International Journal of Interactive Communication Systems and Technologies*, 10(1):86–91.

Hauer, D. (2016). Design of a Smart Contract Based Social Capital Currency System on Ethereum. Bachelor's thesis, *Technical University of Munich*, Germany. Supervised by Sebastian Schams and Georg Groh.

Hauffa, J., Koster, B., Hartl, F., Köllhofer, V., and Groh, G. (2016). Mining twitter for an explanatory model of social influence. In *SocInf@ IJCAI*, pages 3–14.

Hauser, M. (1996). *The evolution of communication*. MIT press.

Hazlitt, H. (1979). *Economics in One Lesson*. Economics (Arlington House). Crown.

Heiberger, R. and Holland, B. (2004). *Statistical Analysis and Data Display: An Intermediate Course with Examples in S-PLUS, R, and SAS*. Springer Texts in Statistics. Springer.

Hevner, A., March, S., Park, J., and Ram, S. (2004). Design science in information systems research. *MIS Q.*, 28(1):75–105.

Hirsch, J. E. (2005). An index to quantify an individual's scientific research output. *Proc Natl Acad Sci U S A*, 102(46):16569–16572.

Hirtle, N. (2018). Social Capital Market Simulation in consideration of Macroeconomic Concepts. Bachelor's thesis, *Technical University of Munich*, Germany. Supervised by Sebastian Schams and Georg Groh.

Ho, T. K. (1995). Random decision forests. In *Document analysis and recognition, 1995., proceedings of the third international conference on*, volume 1, pages 278–282. IEEE.

Höfer, C. (2017). Designing an experiment for the evaluation of social capital assessment and market concepts. Master's thesis, *Technical University of Munich*, Germany. Supervised by Sebastian Schams and Georg Groh.

Hofer, M. and Aubert, V. (2013). Perceived bridging and bonding social capital on twitter: Differentiating between followers and followees. *Comput. Hum. Behav.*, 29(6):2134–2142.

Hoffman, K., Zage, D., and Nita-Rotaru, C. (2009). A survey of attack and defense techniques for reputation systems. *ACM Computer Survey*, 42(1):1–31.

Hoffman, M., Bach, F., and Blei, D. (2010). Online learning for Latent Dirichlet Allocation. In *Advances in neural information processing systems*, pages 856–864.

Hong, S. and Nadler, D. (2012). Which candidates do the public discuss online in an election campaign?: The use of social media by 2012 presidential candidates and its impact on candidate salience. *Government Information Quarterly*, 29(4):455–461.

Investorpedia (2018). The voting rights of shareholders. *https://www.investopedia.com/terms/v/votingright.asp* (accessed May 25, 2018).

Iqbal, R., Murad, M., Mustapha, A., and Sharef, N. (2013). An analysis of ontology engineering methodologies: A literature review. *Research journal of applied sciences, engineering and technology*, 6(16):2993–3000.

Jain, A. K., Murty, M. N., and Flynn, P. J. (1999). Data clustering: a review. *ACM computing surveys (CSUR)*, 31(3):264–323.

Java, A., Song, X., Finin, T., and Tseng, B. (2007). Why we twitter: Understanding microblogging usage and communities. In *Proceedings of the 9th WebKDD and 1st SNA-KDD 2007 Workshop on Web Mining and Social Network Analysis*, WebKDD/SNA-KDD '07, pages 56–65, New York, NY, USA. ACM.

John, O. P. and Srivastava, S. (1999). The big five trait taxonomy: History, measurement, and theoretical perspectives. *Handbook of personality: Theory and research*, 2(1999):102–138.

Jones, S. and Shah, P. (2015). Diagnosing the locus of trust: A temporal perspective for trustor, trustee, and dyadic influences on perceived trustworthiness. *Journal of Applied Psychology*.

Judge, T., Higgins, C., Thoresen, C., and Barrick, M. (1999). The big five personality traits, general mental ability, and career success across the life span. *Personnel psychology*, 52(3):621–652.

Kamvar, S. D., Schlosser, M. T., and Garcia-Molina, H. (2003). The eigentrust algorithm for reputation management in P2P networks. In *Proceedings of the 12th international conference on World Wide Web*, pages 640–651. ACM.

Kas, M., Carley, K. M., and Carley, L. R. (2012). Trends in science networks: understanding structures and statistics of scientific networks. *Social Network Analysis and Mining*, 2(2):169–187.

Keller, E. and Berry, J. (2003). *The Influentials: One American in Ten Tells the Other Nine How to Vote, Where to Eat, and What to Buy*. Free Press.

Kleinberg, J. M. (1999). Authoritative sources in a hyperlinked environment. *Journal of the ACM (JACM)*, 46(5):604–632.

Klügl, F. and Bazzan, A. L. (2012). Agent-based modeling and simulation. *AI Magazine*, 33(3):29.

Kosba, A., Miller, A., Shi, E., Wen, Z., and Papamanthou, C. (2016). Hawk: The blockchain model of cryptography and privacy-preserving smart contracts. In *Security and Privacy (SP), 2016 IEEE Symposium on*, pages 839–858. IEEE.

Koschützki, D., Lehmann, K. A., Peeters, L., Richter, S., Tenfelde-Podehl, D., and Zlotowski, O. (2005). Centrality indices. In *Network analysis*, pages 16–61. Springer.

Krackhardt, D. (1990). Assessing the political landscape: Structure, cognition, and power in organizations. *Administrative science quarterly*, pages 342–369.

Krogerus, M. and Grassegger, H. (2016). Ich habe nur gezeigt, dass es die Bombe gibt. *Das Magazin*, (48-3).

Kywe, S. M., Hoang, T.-A., Lim, E.-P., and Zhu, F. (2012). On Recommending Hashtags in Twitter Networks. *Proceedings of the 4th international conference on Social Informatics*, pages 337–350.

Lampe, C., Johnston, E., and Resnick, P. (2007). Follow the Reader: Filtering Comments on Slashdot. *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 1253–1262.

Lazarsfeld, P., Merton, R., et al. (1954). Friendship as a social process: A substantive and methodological analysis. *Freedom and control in modern society*, 18(1):18–66.

Lee, L. (1999). Measures of distributional similarity. In *Proceedings of the 37th annual meeting of the Association for Computational Linguistics on Computational Linguistics*, pages 25–32. Association for Computational Linguistics.

Lenhart, A., Purcell, K., Smith, A., and Zickuhr, K. (2010). Social media & mobile internet use among teens and young adults. millennials. *Pew internet & American life project*.

Levenshtein, V. (1966). Binary codes capable of correcting deletions, insertions, and reversals. In *Soviet physics doklady*, volume 10, pages 707–710.

Li, N. and Gillet, D. (2013). Identifying influential scholars in academic social media platforms. *Proceedings of the 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining - ASONAM '13*, pages 608–614.

Lin, N. (2002). *Social Capital: A Theory of Social Structure and Action*. Cambridge University Press.

Locke, E. A. (1996). Motivation through conscious goal setting. *Applied and preventive psychology*, 5(2):117–124.

Lu, J., Tang, C., Li, X., and Wu, Q. (2017). Designing socially-optimal rating protocols for crowdsourcing contest dilemma. *IEEE Trans. Information Forensics and Security*, 12(6):1330–1344.

Lu, X. (2012). The relationship of lexical richness to the quality of ESL learners' oral narratives. *The Modern Language Journal*, 96(2):190–208.

Mc Laughlin, H. (1969). SMOG grading – a new readability formula. *Journal of reading*, 12(8):639–646.

Mcknight, D. H. and Chervany, N. L. (1996). The meanings of trust. *Measurement*, 55455(612):86.

Mehrabian, A. (2017). *Nonverbal Communication*. Taylor & Francis.

Mendes, P. N., Jakob, M., and Bizer, C. (2012). DBpedia: A Multilingual Cross-domain Knowledge Base. *LREC*, pages 1813–1817.

Messing, S. and Westwood, S. J. (2014). Selective exposure in the age of social media: Endorsements trump partisan source affiliation when selecting news online. *Communication Research*, 41(8):1042–1063.

Michelson, M. and Macskassy, S. A. (2010). Discovering users' topics of interest on twitter: a first look. In *Proceedings of the fourth workshop on Analytics for noisy unstructured text data*, pages 73–80. ACM.

Miritello, G., Lara, R., Cebrian, M., and Moro, E. (2013). Limited communication capacity unveils strategies for human interaction. *Scientific Reports*, 3(1950).

Mishkin, F. S. and Serletis, A. (2011). *The economics of money, banking, and financial markets*. Pearson Canada.

Molleman, E., Pruyn, J., and Knippenberg, A. (1986). Social comparison processes among cancer patients. *British Journal of Social Psychology*, 25(1):1–13.

Mycoo, M. (2006). Sustainable tourism using regulations, market mechanisms and green certification: A case study of barbados. *Journal of Sustainable Tourism*, 14(5):489–511.

Nahapiet, J. and Ghoshal, S. (2000). Social capital, intellectual capital, and the organizational advantage. In *Knowledge and social capital*, pages 119–157. Elsevier.

Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system. *White paper*.

Navarro, G. (2001). A guided tour to approximate string matching. *ACM computing surveys (CSUR)*, 33(1):31–88.

Neighbors, C. and Knee, C. R. (2003). Self-determination and the consequences of social comparison. *Journal of Research in Personality*, 37(6):529–546.

Nesi, J. and Prinstein, M. (2015). Using social media for social comparison and feedback-seeking: gender and popularity moderate associations with depressive symptoms. *Journal of abnormal child psychology*, 43(8):1427–1438.

Noy, N. F. and McGuinness, D. L. (2001). Ontology development 101: A guide to creating your first ontology. Technical report. *https://protege.stanford.edu/publications/ontology_development/ontology101.pdf* (accessed: March 19, 2018).

NY-Times (2018). Facebook Says Cambridge Analytica Harvested Data of Up to 87 Million Users. *https://www.nytimes.com/2018/04/04/technology/mark-zuckerberg-testify-congress.html* (accessed May 9, 2018).

Ohlberg, M., Shazeda, A., and B, L. (2018). Zentrale Planung, Lokale Experimente. Die komplexe Umsetzung von Chinas gesellschaftlichem Bonitätssystem. *MERICS China Monitor 43*.

Owens, B. D. (2017). Social media is here..."like" it or not. *The American Journal of Sports Medicine*, 45:21–22.

Paavola, J. and Jalonen, H. (2015). An approach to detect and analyze the impact of biased information sources in the social media. In *ECCWS2015-Proceedings of the 14th European Conference on Cyber Warfare and Security 2015: ECCWS 2015*, page 213. Academic Conferences Limited.

Page, L., Brin, S., Motwani, R., and Winograd, T. (1999). The pagerank citation ranking: Bringing order to the web. Technical report, Stanford InfoLab.

Palmer, M. (2000). Records management and accountability versus corruption, fraud and maladministration. *Records Management Journal*, 10(2):61–72.

Papadopoulos, S., Kompatsiaris, Y., Vakali, A., and Spyridonos, P. (2012). Community detection in social media. *Data Mining and Knowledge Discovery*, 24(3):515–554.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al. (2011). Scikit-learn: Machine learning in python. *Journal of machine learning research*, 12(Oct):2825–2830.

Persily, N. (2017). Can democracy survive the internet? *Journal of democracy*, 28(2):63–76.

Petrović, S., Osborne, M., and Lavrenko, V. (2010). The Edinburgh Twitter corpus. In *Proceedings of the NAACL HLT 2010 Workshop on Computational Linguistics in a World of Social Media*, pages 25–26.

Posada, M., Hernandez, C., and Lopez-Paredes, A. (2006). Learning in continuous double auction market. In *Artificial Economics*, pages 41–51. Springer.

Putnam, R. D. (1995). Bowling alone: America's declining social capital. *Journal of democracy*, 6:65–65.

Quinlan, J. R. (1986). Induction of decision trees. *Machine learning*, 1(1):81–106.

Quinlan, J. R. (1993). *C4.5: Programs for Machine Learning*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA.

R. Botsman (2017). Big data meets Big Brother as China moves to rate its citizens. *http://www.wired.co.uk/article/chinese-government-social-credit-score-privacy-invasion* (accessed May 20, 2018).

R. Verma, K. Singh, A. G. (2010). Article: Simulation: An effective marketing tool. *International Journal of Computer Applications*, 4(11):8–12.

Rainie, H., Anderson, J. Q., and Albright, J. (2017). *The future of free speech, trolls, anonymity and fake news online*. Pew Research Center Washington, DC.

Ramage, D., Dumais, S. T., and Liebling, D. J. (2010). Characterizing microblogs with topic models. *ICWSM*, 10(1):16.

Rao, A., Spasojevic, N., Li, Z., and DSouza, T. (2015). Klout score: Measuring influence across multiple social networks. In *Big Data (Big Data), 2015 IEEE International Conference on*, pages 2282–2289. IEEE.

Raschka, S. (2015). *Python Machine Learning*. Packt Publishing.

Rastogi, T. (2016). A power law approach to estimating fake social network accounts. *arXiv preprint arXiv:1605.07984*.

Razzaghzadeh, S., Navin, A. H., Rahmani, A. M., and Hosseinzadeh, M. (2017). Probabilistic modeling to achieve load balancing in expert clouds. *Ad Hoc Networks*, 59:12–23.

Recuero, R., Araujo, R., and Zago, G. (2011). How does social capital affect retweets? In *ICWSM*.

Resnick, P., Zeckhauser, R., Friedman, E., and Kuwabara, K. (2000). Reputation Systems: Facilitating Trust in Internet Interactions. *Communications of the ACM*, 43(12):45–48.

Revithi, P. (2017). Design of Social Capital Markets. Guided research project, *Technical University of Munich*, Germany. Supervised by Sebastian Schams and Georg Groh.

Richterich, A. (2013). "Karma, Precious Karma!" Reddit and the Front Page's Econometrisation. *Journal of Peer Production*, (May):1–12.

Robison, L., Schmid, A., and Siles, M. (2002). Is social capital really capital? *Review of Social Economy*, 60(1):1–21.

Robles, P. (2011). Twitter isn't very social: study. *Econsultancy*.

Rowe, R., Creamer, G., Hershkop, S., and Stolfo, S. (2007). Automated social hierarchy detection through email network analysis. *Joint 9th WebKDD and 1st SNA-KDD Workshop 2007 on Web Mining and Social Network Analysis. Held in conjunction with 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD 2007*, pages 109–117.

Runkler, T. (2015). *Data Mining: Modelle und Algorithmen intelligenter Datenanalyse*. Computational Intelligence. Springer Fachmedien Wiesbaden.

Sang, J. A., Ismail, R., and Boyd, C. (2007). A survey of trust and reputation systems for online service provision. *Decision Support Systems; Emerging Issues in Collaborative Commerce*, 43(2):618–644.

Schaeffer, S. E. (2007). Graph clustering. *Computer science review*, 1(1):27–64.

Schmidt, M. (2017). Social Capital Market Simulation. Bachelor's thesis, *Technical University of Munich*, Germany. Supervised by Sebastian Schams and Georg Groh.

Schmitt, D., Realo, A., Voracek, M., and Allik, J. (2008). Why can't a man be more like a woman? Sex differences in Big Five personality traits across 55 cultures. *Journal of personality and social psychology*, 94(1):168.

Schütze, H., Manning, C. D., and Raghavan, P. (2008). *Introduction to information retrieval*, volume 39. Cambridge University Press.

Shapiro, S. and Wilk, M. (1965). An analysis of variance test for normality (complete samples). *Biometrika*, 52(3/4):591–611.

Sheldon, K., Ryan, R., Rawsthorne, L., and Ilardi, B. (1997). Trait self and true self: Cross-role variation in the big-five personality traits and its relations with psychological authenticity and subjective well-being. *Journal of personality and social psychology*, 73(6):1380.

Sidani, J., Shensa, A., Radovic, A., Miller, E., Colditz, J., Hoffman, B., Giles, L., Primack, B., et al. (2016). Association between social media use and depression among US young adults. *Depression and anxiety*, 33(4):323–331.

Simon, H. A. (1955). On a class of skew distribution functions. *Biometrika*, 42(3/4):425–440.

Skopik, F., Truong, H. L., and Dustdar, S. (2009). Trust and reputation mining in professional virtual communities. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 5648 LNCS:76–90.

Spees, K. and Lave, L. B. (2007). Demand response and electricity market efficiency. *The Electricity Journal*, 20(3):69–85.

Stark, S. (2001). *Sichere IT-Kommunikation über unsichere Netze*. Diplom.de Verlag.

Statista (2017). Distribution of Facebook users worldwide as of January 2017, by age and gender. *https://www.statista.com/statistics/376128/facebook-global-user-age-distribution/* (accessed January 28, 2018).

Stevenson, A. (2010). *Oxford Dictionary of English*. Oxford Dictionary of English. OUP Oxford.

Steyvers, M. and Griffiths, T. (2007). Probabilistic topic models. *Handbook of latent semantic analysis*, 427(7):424–440.

Studer, R., Benjamins, R., and Fensel, D. (1998). Knowledge engineering: Principles and methods. *Data & Knowledge Engineering*, 25(1-2):161–198.

Su, H., Tang, J., and Hong, W. (2012). Learning to diversify expert finding with subtopics. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, pages 330–341. Springer.

Sun, J. (2011). *A survey of models and algorithms for social influence analysis*, volume 54.

Süß, P. (2018). Evaluation of Social Capital Markets and Discussion of Use-Cases. Bachelor's thesis, *Technical University of Munich*, Germany. Supervised by Sebastian Schams and Georg Groh.

Tang, J., Zhang, J., Yao, L., Li, J., Zhang, L., and Su, Z. (2008). ArnetMiner: Extraction and Mining of Academic Social Networks. *Proceedings of the 14th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 990–998.

Tapscott, D. and Tapscott, A. (2016). *Blockchain Revolution: How the Technology Behind Bitcoin Is Changing Money, Business, and the World*. Penguin Publishing Group.

TheTelegraph (2017). China's 'social credit' system bans millions from traveling. *https://www.telegraph.co.uk/news/2018/03/24/chinas-social-credit-system-bans-millions-travelling/* (accessed May 20, 2018).

Thorndike, E. L. (1920). A constant error in psychological ratings. *Journal of applied psychology*, 4(1):25–29.

Toader, S., Morar, L., Câmpean, E., et al. (2015). A comparative performance analysis of the credit bureau of romania & schufa holding ag in germany. *Theoretical and Applied Economics*, 22(Special (II)):331–342.

True, C. (2017). The best medium for relationships. *ESSAI*, 15(1):35.

Ullrich, M. (2014). Casual Influence-Structures on Facebook. Bachelor's thesis, *Technical University of Munich*, Germany. Supervised by Jan Hauffa and Georg Groh.

Underwood, J. D., Kerlin, L., and Farrington-Flint, L. (2011). The lies we tell and what they say about us: Using behavioural characteristics to explain facebook activity. *Computers in Human Behavior*, 27(5):1621–1626.

U.S. Department of the Treasury (2017). Distribution of currency and coins. *https://www.treasury.gov/about/education/Pages/distribution.aspx* (accessed February 12, 2018).

Valenzuela, S., Park, N., and Kee, K. F. (2009). Is there social capital in a social network site?: Facebook use and college student's life satisfaction, trust, and participation. *Journal of Computer-Mediated Communication*, 14(4):875–901.

Van Hout, R. and Vermeer, A. (2007). Comparing measures of lexical richness. *Modelling and assessing vocabulary knowledge*, pages 93–115.

Van Vlasselaer, V., Bravo, C., Caelen, O., Eliassi-Rad, T., Akoglu, L., Snoeck, M., and Baesens, B. (2015). Apate: A novel approach for automated credit card transaction fraud detection using network-based extensions. *Decision Support Systems*, 75:38–48.

Varshney, M. (2017). Models of Topic-Specic Reputation / Social Capital in Online Social Networks. Master's thesis, *Technical University of Munich*, Germany. Supervised by Sebastian Schams and Georg Groh.

Wang, B. and Zhuang, J. (2018). Rumor response, debunking response, and decision makings of misinformed twitter users during disasters. *Natural Hazards*, pages 1–18.

Wargo, E. (2006). How many seconds to a first impression? *APS Observer*, 19(7).

Watts, D. J. and Strogatz, S. H. (1998). Collective dynamics of small-world networks. *nature*, 393(6684):440.

Weng, J., Lim, E.-P., Jiang, J., and He, Q. (2010). Twitterrank: Finding topic-sensitive influential Twitterers. *Proceedings of the Third ACM International Conference on Web Search and Data Mining*, pages 261–270.

Weninger, T., Zhu, X., and Han, J. (2013). An exploration of discussion threads in social news sites: A case study of the reddit community. In *Proceedings of the 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining*, ASONAM '13, pages 579–583, New York, NY, USA. ACM.

Wolfers, J. and Zitzewitz, E. (2004). Prediction markets. *Journal of economic perspectives*, 18(2):107–126.

Wu, X., Kumar, V., Quinlan, R., Ghosh, J., Yang, Q., Motoda, H., McLachlan, G., Ng, A., Liu, B., Yu, P., Zhou, Z., Steinbach, M., Hand, D., and Steinberg, D. (2007). Top 10 algorithms in data mining. *Knowl. Inf. Syst.*, 14(1):1–37.

Yan, X. (2009). *Linear Regression Analysis: Theory and Computing*. World Scientific Publishing Company Pte Limited.

Yang, C., Harkreader, R. C., and Gu, G. (2011). Die free or live hard? empirical evaluation and new design for fighting evolving twitter spammers. In *International Workshop on Recent Advances in Intrusion Detection*, pages 318–337. Springer.

Yang, Z., Hong, L., and Davison, B. D. (2010). Topic-driven multi-type citation network analysis. *Adaptivity, Personalization and Fusion of ...*, pages 24–31.

Zettler, S. (2016). Evaluation of Concepts for Social Capital Integration in Social Networks. Bachelor's thesis, *Technical University of Munich*, Germany. Supervised by Sebastian Schams and Georg Groh.

Zeynalov, R. (2018). Cross-medial Machine Learning based Social Capital Analysis. Master's thesis, *Technical University of Munich*, Germany. Supervised by Sebastian Schams and Georg Groh.

Zhang, J., Arbor, A., Ackerman, M. S., and Arbor, A. (2005). Searching For Expertise in Social Networks: A Simulation of Potential Strategies. *GROUP '05 Proceedings of the 2005 international ACM SIGGROUP conference on Supporting group work*, pages 71–80.

Zhao, H. and Li, X. (2009). H-trust: a group trust management system for peer-to-peer desktop grid. *Journal of Computer Science and Technology*, 24(5):833–843.

Zhou, R. and Hwang, K. (2007). Powertrust: A robust and scalable reputation system for trusted peer-to-peer computing. *IEEE Transactions on parallel and distributed systems*, 18(4):460–473.

Ziegler, C.-N. (2009). *On Propagating Interpersonal Trust in Social Networks*, pages 133–168. Springer London, London.