

TECHNISCHE UNIVERSITÄT MÜNCHEN

Fakultät für Mathematik — Lehrstuhl M12 (Mathematische Modellierung
biologischer Systeme)

**Quantitative mathematical modeling of
biological processes using reduced
parameter and data spaces**

Anna Kristina Fiedler

Vollständiger Abdruck der von der Fakultät Mathematik der Technischen Uni-
versität München zur Erlangung des akademischen Grades eines

Doktors der Naturwissenschaften (Dr. rer. nat.)

genehmigten Dissertation.

Vorsitzender:

Prof. Dr. Oliver Junge

Prüfer der Dissertation:

1. Prof. Dr. Jan Hasenauer
2. Prof. Dr. Barbara Kaltenbacher
3. Prof. Dr. Christina Kuttler

Die Dissertation wurde am 27.05.2020 bei der Technischen Universität München
eingereicht und durch die Fakultät für Mathematik am 10.11.2020 angenommen.

Acknowledgement

First of all, I want to thank my thesis supervisor Jan Hasenauer for his scientific guidance and input, for his general support, his feedback and his patience. I further want to thank Fabian Theis for his guidance, for giving me the opportunity to join the Institute of Computational Biology and for fostering an inspiring scientific and working environment.

I want to thank my thesis committee, Barbara Kaltenbacher for the scientific discussions and for hosting me at the Alpen Adria University in Klagenfurt, Christina Kuttler for her feedback and taking the time for assessing my thesis and Oliver Junge for chairing the examination. I am further grateful to the TUM Graduate School and the Global Challenges for Women in MathScience Program for providing financial support for my research stays abroad.

I was very happy to have the possibility to be part of two amazing working environments, the Zentrum Mathematik at the TUM and the Institute of Computational Biology at the Helmholtz Zentrum München and I want to thank all my former colleagues at both places for making my time there unforgettable. I would like to thank Donna Ankerst, Hannes Petermeier and Johannes Müller for their feedback and their support regarding teaching matters. I want to thank Andreas Strobl, Atefeh Kazeroonian Bendix Koopmann, Benjamin Ballnus, Carolin Loos, Dantong Wang, Daniel Weindl, Dennis Rickert, Elba Raimúndez-Álvarez, Erika Dudkin, Eva Stadler, Fabian Fröhlich, Jakob Vanhoefer, Johanna Tolksdorf, Josef Höfler, Karin Tichmann, Karolina Worf, Katharina Selig, Leonard Schmiester, Linda Krause, Lisa Bast, Maren Büttner, Michael Strobel, Norbert Krautenbacher, Paul Stapor, Polina Lakrisenko, Sabine Hug, Sabrina Hross, Simon Merkt, Sonja Grill, Stephanie Troppmann, Yannik Schälte and Yiyao Chen for being my office mates, colleagues, thesis writer club mates, my proof readers, for the scientific and unscientific discussions, for being my lunch and coffee break company or any subset of the above.

Next, I would like to thank my collaborators, especially Romana Boiger and David Fischer, for the fruitful cooperations and also Vitaly Ganusov for hosting me at the University of Tennessee in Knoxville and providing a new scientific perspective.

Lastly, I want to thank my friends and especially my parents, Marlies und Michael, and my siblings Friederike and Sebastian for supporting and encouraging me all throughout my studies, my PhD studies and the writing process. Moritz, thank you for everything! This thesis would not have been finished without you. Thank you!

Abstract

Quantitative mathematical models are widely used to describe biological processes and analyze biological data. These models usually contain unknown parameters that have to be estimated from experimental data. To estimate the model parameters, optimization problems have to be solved, which usually entails repeated model simulation. The difficulty of model simulation and parameter estimation usually increases for problems with a large number of unknown model parameters, a large state space or high dimensional experimental data. The project discussed in this thesis investigated how to use reduced parameter and data spaces to facilitate robust and efficient modeling and parameter estimation.

First, we considered ordinary and partial differential equation models with the additional constraint that the model is initially in steady state. In this setting, a pair of model parameters and initial conditions that satisfy the steady-state constraint has to be inferred from the data. We used the geometry of the steady-state manifold, i.e., the manifold of model parameters and states that fulfill the steady-state constraint, to reduce the space of optimization variables by choosing a state update on the steady-state manifold for a given parameter update. We developed two methods for optimization, an approach based on retraction mappings and a continuous analogue approach. Further, we proved stability of optimal parameter-state pairs for the continuous analogue. We further investigated the influence of a hyper parameter of the continuous analogue approach, i.e., the retraction factor, on convergence and runtime and proposed a lower bound. Lastly, we applied the proposed approaches to simulated and real data and compared them to established methods for constrained optimization. Here, the new approaches demonstrated better convergence and efficiency than the established methods.

In the second part, we consider the modeling of heterogeneous cell populations observed using high-dimensional single-cell measurements. Instead of modeling the process dynamics in the high-dimensional space, we used a reduction of the high-dimensional data space to a one-dimensional trajectory to construct a diffusion-advection-reaction model of the dynamics of heterogeneous cell populations. The model describes the proliferation dynamics, i.e., cell division and death, and the differentiation dynamics along the extracted

one dimensional cell state trajectory. To describe differentiation processes with two different possible cell fates, we extended the diffusion-advection-reaction model to a branching model on two branches. Further, we developed a statistical framework for parameter estimation that combines different data types. We applied the developed model and parameter estimation framework to single-cell RNA sequencing data for T-cell maturation in thymus and estimated proliferation and drift rates that give some unbiased indication of the properties of the maturation process.

Overall, this thesis presents two novel approaches for high-dimensional problems: inference problems as well as data sets. The contributions might help to facilitate the analysis of biological processes and data that was previously limited by the problem size.

Zusammenfassung

Quantitative mathematische Modelle werden häufig verwendet, um biologische Prozesse zu beschreiben und biologische Daten zu analysieren. Diese Modelle enthalten normalerweise unbekannte Parameter, die aus experimentellen Daten geschätzt werden müssen. Um die Modellparameter zu schätzen, müssen Optimierungsprobleme gelöst werden, die in der Regel eine wiederholte Modellsimulation erfordern. Die Schwierigkeit von Modellsimulation und Parameterschätzung steigt in der Regel mit der Anzahl unbekannter Modellparameter bzw. mit der Dimensionalität der experimentellen Daten. In dieser Arbeit wurde untersucht, wie reduzierte Parameter- und Datenräume verwendet werden können, um eine robuste und effiziente Modellierung und Parameterschätzung zu ermöglichen.

Im ersten Teil betrachten wir zunächst gewöhnliche und partielle Differentialgleichungsmodelle mit der zusätzlichen Nebenbedingung, dass der Anfangszustand ein stationärer Punkt ist. Für diese Probleme müssen sowohl die Modellparameter als auch die Anfangsbedingungen aus den Daten geschätzt werden und dabei die Stationaritätsbedingung erfüllen. Wir verwendeten die Geometrie der Mannigfaltigkeit von Modellparametern und Anfangszuständen, auf der die Stationaritätsbedingung erfüllt ist, um den Raum der Optimierungsvariablen zu verringern, indem wir für eine gegebene Parameteraktualisierung eine Zustandsaktualisierung auf der Mannigfaltigkeit wählen. Wir entwickelten zwei Optimierungsmethoden, einen Ansatz, der auf Retraktionsmappings basiert und einem Ansatz der ein kontinuierliches Analogon zu diskreten Optimierungsmethoden verwendet. Für das kontinuierliche Analogon bewiesen wir Stabilität von optimalen Parameterzustandspaaren. Wir untersuchten weiter den Einfluss eines Hyperparameters des kontinuierlichen Analogons, des Retraktionsfaktors, auf Konvergenz und Laufzeit und schlugen eine Untergrenze vor. Zuletzt wendeten wir die vorgeschlagenen Methoden auf simulierte und reale Daten an und verglichen sie mit etablierten Methoden zur eingeschränkten Optimierung. Hier zeigten die neuen Ansätze eine bessere Konvergenz und Effizienz als die etablierten Methoden.

Im zweiten Teil betrachten wir die Modellierung heterogener Zellpopulationen, die mit hochdimensionalen Einzelzellmessungen beobachtet wurden. Anstatt die Prozessdynamik

im hochdimensionalen Raum zu modellieren, verwendeten wir eine Reduktion des hochdimensionalen Datenraums auf eine eindimensionale Trajektorie, um ein Diffusions-Advektions-Reaktionsmodell der Dynamik heterogener Zellpopulationen zu erstellen. Das Modell beschreibt die Proliferationsdynamik, d.h. Zellteilung und Tod, und die Differenzierungsdynamik entlang der extrahierten eindimensionalen Zellzustandstrajektorie. Um Differenzierungsprozesse mit zwei unterschiedlichen möglichen Endstadien zu beschreiben, erweiterten wir das Diffusions-Advektions-Reaktions-Modell auf zwei Zweige zu einem Verzweigungsmodell. Darüber hinaus entwickelten wir einen statistischen Rahmen für die Parameterschätzung, das verschiedene Datentypen kombiniert. Wir wendeten das entwickelte Modell- und Parameterschätzungsgerüst auf Einzelzell-RNA-Sequenzierungsdaten für die T-Zell-Reifung im Thymus an und schätzen Proliferations- und Driftraten, die einen unvoreingenommenen Hinweis auf die Eigenschaften des Reifungsprozesses geben.

Insgesamt werden in dieser Arbeit zwei neuartige Ansätze für hochdimensionale Probleme vorgestellt: Inferenzprobleme sowie Datensätze. Die Methoden könnten dazu beitragen, die Analyse biologischer Prozesse und Daten zu ermöglichen, die zuvor durch die Problemgröße limitiert war.

Contents

1. Introduction	1
1.1. Steady-state constraints for ODE and PDE constrained optimization	2
1.2. Population dynamics based on high dimensional single-cell measurements .	4
1.3. Contribution of this thesis	5
1.4. Outline of this thesis	7
2. Background	9
2.1. Dynamical systems	9
2.1.1. Ordinary differential equations	10
2.1.2. Partial differential equations	11
2.1.3. Asymptotic behavior	13
2.2. Statistical models and uncertainty analysis	17
2.2.1. Maximum likelihood estimation	17
2.2.2. Uncertainty analysis and confidence intervals	19
2.2.3. Regularization and Bayesian inference	20
2.2.4. Model selection	21
2.3. Inverse problems and ODE/PDE constrained optimization	21
2.3.1. Global optimization	23
2.3.2. Local optimization	23
2.3.3. Continuous analogues	24
2.3.4. Gradients and sensitivities	24
3. ODE constrained problems with additional steady-state constraints	27
3.1. Introduction and mathematical problem formulation	27
3.1.1. Formulation of mathematical model	28
3.1.2. Parameter estimation problem	30
3.1.3. Established optimization approaches	31
3.1.4. Problem formulation	32
3.2. Tailored Algorithms	32
3.2.1. Manifold of steady states	32

Contents

3.2.2.	Hybrid optimization method	34
3.2.3.	Continuous analogues for optimization	35
3.2.4.	Asymptotic stability of the optimal parameter-state pair	37
3.3.	Application and evaluation	41
3.3.1.	Conversion process	41
3.3.2.	Behavior in the presence of bistability, bifurcations and oscillations .	44
3.3.3.	NGF-induced ERK signaling in primary sensory neurons	48
3.3.4.	Raf/MEK/ERK signaling in HeLa cells after release from S-phase arrest	54
3.4.	Discussion	62
4.	PDE constrained problems with additional steady-state constraints	65
4.1.	Introduction and mathematical problem formulation	65
4.1.1.	Formulation of mathematical model	66
4.1.2.	Parameter estimation problem	67
4.1.3.	Problem formulation	68
4.2.	Continuous analogue	69
4.2.1.	Elliptic PDE constraints	69
4.2.2.	Elliptic and parabolic PDE constraints	71
4.3.	Local stability and convergence to a local optimum	72
4.3.1.	Elliptic PDE constraints	74
4.3.2.	Elliptic and parabolic PDE constraints	79
4.4.	Application	80
4.4.1.	Model formulation	80
4.4.2.	Numerical implementation	82
4.4.3.	Simulated data	82
4.4.4.	Optimization	83
4.4.5.	Comparison of continuous analogue and discrete iterative procedure	85
4.4.6.	Evaluation of retraction factor influence	85
4.5.	Discussion	86
5.	Modeling of population dynamics on reduced data spaces	89
5.1.	Introduction and problem formulation	89
5.1.1.	Biological background	90
5.1.2.	Data structure	92
5.1.3.	Problem formulation	92
5.2.	Mathematical model	94
5.2.1.	Non-branching model	94

5.2.2.	Branching model	96
5.2.3.	Numerical scheme for simulation	97
5.3.	Parameter estimation	100
5.3.1.	Parameterization using splines	100
5.3.2.	Likelihood	101
5.3.3.	Regularization	105
5.3.4.	Selection of the regularization hyper-parameter	106
5.4.	Application to T-cell maturation data	107
5.4.1.	Experiment and data processing	107
5.4.2.	Numerical implementation	108
5.4.3.	The K-S distance is overly sensitive for the T-cell data	110
5.4.4.	Model fit and regularization	110
5.4.5.	Parameter estimates identify checkpoints and selection in T-cell maturation	112
5.5.	Discussion	112
6.	Conclusion and outlook	115
	Appendix A. Analytical steady state for the Raf/MEK/ERK signaling model	119
	Appendix B. Finite volume discretization for the branching model	123

Chapter 1.

Introduction

Mathematical models are valuable tools for studying biological processes. These mathematical formulations permit the incorporation of prior knowledge about the underlying processes, data integration and hypothesis testing (Kitano, 2002; Klipp et al., 2005). Mechanistic mathematical modeling facilitates a deeper understanding of the relevant mechanisms (Kitano, 2002), the derivation of biological hypotheses, e.g., predictions of future or unobserved states or the response to perturbations, and experimental design for hypothesis testing. Based on the results of experiments, models can in turn be adapted and refined (Kitano, 2002). Once a model captures the biological process of interest in necessary detail, experiments can be performed *in silico* thereby saving resources (see e.g., Fröhlich et al. (2018)). In this work, we focus on ordinary differential equation (ODE) and partial differential equation (PDE) models, two model classes which are widely used in mathematical and systems biology. ODEs and PDEs provide a mechanistic and quantitative description of the temporal evolution of the states of a biological system and are used to model various processes (Klipp et al., 2005; Perthame, 2015).

As the descriptive power of a mathematical model depends on the accuracy of parameters, the model parameters have to be estimated from experimental data (Isakov, 2006; Tarantola, 2005). The inverse problem of parameter estimation from biological experiments is commonly challenging due to high noise, incomplete observations and nonlinearity. In the biological context, maximum likelihood approaches are often used for the parameter estimation. For ODE or PDE models this leads to ODE or PDE constrained optimization problems. These optimization problems are usually non-convex, non-linear and computationally demanding, as the models usually have to be simulated in every optimization step. Hence, efficient and robust optimization methods are required (Banga, 2008; Chou and Voit, 2009). For ODE and PDE constrained optimization, established methods include global optimization (Moles et al., 2003; Villaverde et al., 2015) and multi-start local opti-

mization (Raue et al., 2013b) for ODEs and several numerical approaches for PDEs (Banks and Kunsich, 1989; Bock et al., 2013; Carvalho et al., 2015; Hinze et al., 2009; Ito and Kunisch, 2008; Xun et al., 2013).

Due to technical advancements related to experimental techniques in recent years, it is possible to gather larger amounts of biological measurements at a higher resolution. The depth and dimensionality of the resulting data is steadily increasing, requiring more detailed models. Accordingly, parameter estimation is becoming increasingly time and resource consuming. Model size and complexity also increase to incorporate all available and essential knowledge, e.g., in (Fröhlich et al., 2018; Hasenauer et al., 2015; Holzhütter et al., 2012; Karr et al., 2012). The (i) large number of species and parameters considered in a model and (ii) the high dimension of the data often impede efficient numerical simulation and parameter estimation. A possible answer to this problem is to reduce the dimension of the problem either on the model or the data side.

In this thesis, we investigate how to use manifolds to reduce the dimensionality of optimization problems and efficiently estimate parameters. We follow two different routes:

- We consider problems with steady-state constraints and exploit geometric approaches to reduce the dimensionality of the space of optimization variables.
- We consider high dimensional data sets and use dimension reduction to render them accessible to dynamic modeling.

A brief introduction to the topics is provided in Section 1.1 and 1.2. In Section 1.3, we further give an overview of the contributions of this thesis and a list of publications and in Section 1.4 the structure of the thesis is outlined.

1.1. Steady-state constraints for ODE and PDE constrained optimization

Models of biological processes often contain unknown parameters and initial conditions, that need to be estimated from experimental data. The data can be quantitative information, such as the value of an observable at different time points, or qualitative. Typical qualitative information is that the system is initially in steady state. Steady states occur in many biological applications. For example, due to processes happening on different time scales, a system often remains in a pseudo steady state in relation to the time scale considered (Klipp et al., 2005). A process might also be regulated to stay in homeosta-

1.1. Steady-state constraints for ODE and PDE constrained optimization

sis and hence naturally favor a steady state (e.g., in O’Dea et al. (2007)) or it might be sufficient and economic to only measure the steady state before and after a perturbation (e.g., in Molinelli et al. (2013)).

The information that the system is in steady state constrains the space of feasible parameter-state combinations. Hence, neglecting this information might lead to wrong parameterization and unreliable predictions. The steady-state constraint is incorporated into the parameter estimation by formulating a non-linear equality constraint for the optimization variables, i.e. parameters and state, in the corresponding optimization problem. This additional non-linear equality constraints complicates the estimation (Rosenblatt et al., 2016).

Solving the optimization problem with the non-linear equality constraints can be circumvented by computing steady states analytically as function of the parameters using py-substitution, a method based on computer algebra (Rosenblatt et al., 2016). This approach was initially developed for enzyme-catalyzed reactions (Cornish-Bowden, 1977; King and Altman, 1956) and later extended to include more reaction kinetics (Chou, 1990; Feliu and Wiuf, 2012; Halász et al., 2013; Loriaux et al., 2013) and to ensure positivity of solutions (Rosenblatt et al., 2016). However, this method depends on the presence of sufficient degrees of freedom in the system of ODE constraints and is not applicable to PDEs.

For equality constraints that correspond to matrix manifolds, optimization on matrix manifold techniques have been developed (Absil et al., 2008; Bertsekas, 1999). By restricting the search space to the non-linear manifold of admissible states, these methods often achieve favorable numerical properties. Optimization on matrix manifold algorithms utilize retraction mappings to map from the tangent spaces to points on the manifold and to efficiently construct optimizer updates on the manifold for discrete iterative optimization (Absil et al., 2008). However, these retraction operators have to be constructed specifically for the problems at hand which is usually a not negligible task (Absil et al., 2008).

Apart from discrete iterative optimization algorithms, also continuous analogues are employed for optimization. Continuous analogues are dynamical systems whose stable equilibria are optimal points of the optimization problem. Hence, the solution of a continuous analogue converges to the optimum of the corresponding problem. Continuous analogues have been constructed for gradient descent methods, where they take the form of gradient flows (Wang et al., 2003), and also for Newton-Raphson and Levenberg-Marquardt methods (Kaltenbacher et al., 2002; Tanabe, 1985). Among others, methods based on

continuous analogues have been employed for unconstrained (Botsaris, 1978; Brown and Bartholomew-Biggs, 1989) and constrained optimization (Wang et al., 2003), for solving nonlinear equations (Airapetyan et al., 2000; Tanabe, 1985) and for finding saddle points (Dürr and Ebenbauer, 2011; Kose, 1956). For many problems these methods tend to have larger regions of attraction and converge more robustly compared to their discrete iterative counterparts (Tanabe, 1985).

1.2. Population dynamics based on high dimensional single-cell measurements

The massive advances in the experimental analysis of single cells using, e.g., single-cell RNA sequencing (Klein et al., 2015), mass cytometry (Bandura et al., 2009) and high-dimensional flow cytometry (Saeys et al., 2016) enable the resolution of biological processes on an unprecedented scale (Angerer et al., 2017; Linnarsson and Teichmann, 2016). By measuring high dimensional properties like the amount of individual transcripts or proteins for each single cell in a population of cells, these techniques provide insight into the state, cell type and function of the cells. However, the large number of dimensions in the data set impedes the inference of population dynamics using mathematical models.

To reduce the dimensionality of single-cell data, methods for pseudotemporal ordering can be employed (Haghverdi et al., 2016; Trapnell et al., 2014). As cells often develop asynchronously, the considered cell populations are heterogeneous and the generated snapshots contain cells in a variety of states. Pseudotemporal orderings sort these cells with respect to their similarity and construct a developmental trajectory of cells in the high dimensional space. To obtain a broader picture of the cell state trajectory, snapshots at different time points can be pooled. Additionally, pseudotemporal orderings are applied to identify branching of the trajectory to different cell fates (e.g., in Haghverdi et al. (2016)).

The trajectory recovered by pseudotemporal orderings is inherently static and not linked to the time points at which the experiments were performed. Without this link there is no description of the population dynamics. Additionally, proliferation dynamics, i.e. cell death and division, are not accessible from this static description.

In the literature, population dynamics, which typically consist of differentiation and proliferation, are often addressed using discrete developmental stages and transitions between them resulting in ODE models, e.g., in (Lander et al., 2009). However, discrete developmental stages are usually defined in the data based on marker proteins and the sorting of

cells relies on cut offs and neglects intermediate stages. For a single static snapshot a model with continuous cell state was developed in (Weinreb et al., 2018) to infer trajectories using a population balance analysis approach. As the authors discuss in their manuscript, the identification of dynamics from static snapshots is limited by the ambiguity of dynamics that can lead to the same snapshot. A different approach to model differentiation dynamics is a coupled ODE-PDE model that describes the continuous differentiation between two discrete developmental stages using a transport equation (Doumic et al., 2011). In (Schiebinger et al., 2019) an optimal transport framework was used to model the (time) discrete transitions of the cell density from one time point to the cell density at the next time point by minimizing transport costs. A dynamic differentiation model based on stochastic processes was proposed in (Hashimoto et al., 2016) and applied to relatively low dimensional gene expression data.

1.3. Contribution of this thesis

Based on the two sections above, we formulate the main questions answered in this thesis.

How can the structure of steady-state constraints be exploited to reduce the dimensionality of ODE and PDE constrained optimization problems with additional steady-state constraints and to develop robust and efficient parameter estimation algorithms? (See Problem 3.1.1, Problem 4.1.1 and Problem 4.1.2)

How can cell population dynamics observed in high-dimensional data sets be modeled, and how can the parameters of these processes be inferred? (See Problem 5.1.1 and Problem 5.1.2)

To address these questions, we developed the following approaches:

- We constructed novel approaches for robust and efficient parameter estimation for ODE and PDE constrained optimization with additional steady-state constraints that use the first order geometry of the steady-state manifold. These approaches include ideas from retraction operators and ideas from continuous analogues.
- We proved asymptotic stability of the continuous analogue approach for both the ODE and PDE constrained problems implying convergence to a local optimum.
- We developed an approach to model the dynamics of heterogeneous cell populations using observations generated by high dimensional single-cell measurements.

Chapter 1. Introduction

- We developed a statistical model for parameter estimation for the aforementioned population model from sequencing data and estimated parameters for T-cell differentiation data in thymus.

The aforementioned developments have been published in the following peer-reviewed papers that constitute the basis of this thesis. Parts and results of this thesis will therefore correspond or be identically to the publications. The corresponding paper is additionally mentioned at the beginning of the respective chapter.

- **A. Fiedler**, S. Raeth, F.J. Theis, A. Hausser, J. Hasenauer, Tailored parameter optimization methods for ordinary differential equation models with steady-state constraints. *BMC Systems Biology*, 10(1):80, 2016.
- R. Boiger*, **A. Fiedler***, J. Hasenauer, B. Kaltenbacher, Continuous analogue to iterative optimization for PDE-constrained inverse problems. *Inverse Problems in Science and Engineering*, 27:6, 710–734, 2019.
- D.S. Fischer*, **A.K. Fiedler***, E.M. Kernfeld, R.M.J. Genga, A. Bastidas-Ponce, M. Bakhti, H. Lickert, J. Hasenauer, R. Maehr, F.J. Theis, Inferring population dynamics from single-cell RNA-sequencing time series data. *Nature Biotechnology*, 37(4):461–468, 2019.

In addition to the publications above, I also contributed to further projects which led to the publication of the following papers that are not part of the thesis.

- C. Loos, **A. Fiedler**, J. Hasenauer, Parameter estimation for reaction rate equation constrained mixture models. In E. Bartocci, P. Lio, and N. Paoletti, editors, *Proceedings of the 14th International Conference on Computational Methods in Systems Biology, Lectures in Bioinformatics*, 186–200. Springer International Publishing, 2016.
- S. Hross*, **A. Fiedler***, F.J. Theis, J. Hasenauer, Quantitative comparison of competing PDE models for Pom1p dynamics in fission yeast. In R. Findeisen, E. Bullinger, E. Balsa-Canto, and K. Bernaerts, editors, *Proceedings of 6th IFAC Conference on Foundations of Systems Biology in Engineering*, 264–269. IFAC-PapersOnLine, 2016.

*These authors contributed equally to the paper

- P. Stapor, D. Weindl, B. Ballnus, S. Hug, C. Loos, **A. Fiedler**, S. Krause, S. Hross, F. Fröhlich, J. Hasenauer, PESTO: Parameter ESTimation TOolbox. *Bioinformatics* 34, 705–707, 2018.

1.4. Outline of this thesis

In Chapter 2, we discuss the mathematical concepts and preliminaries relevant to this work. We briefly introduce dynamical systems, in particular ODEs and PDEs, and their asymptotic behavior. We further outline maximum likelihood based parameter estimation and uncertainty analysis. Finally, we provide a brief overview of ODE/PDE constrained optimization and elementary optimization concepts.

In Chapter 3, we introduce ODE constrained optimization problems with steady-state constraints and develop two estimation approaches based on the first order geometry of the dynamical system. We use singular perturbation theory to prove asymptotic stability for the continuous analogue approach. Lastly, we apply the newly developed approaches to simulated and biological data.

In Chapter 4, we introduce the setting of steady-state constraints for PDE constrained optimization problems and extend the continuous analogue approach developed in Chapter 3 to this case. Further, we prove asymptotic stability of the optimum using a Lyapunov function. Subsequently, we apply the novel approach to simulated data.

In Chapter 5, we present a PDE model for cell differentiation along a cell state axis, develop a likelihood for parameter estimation and apply the model to data of T-cell maturation in thymus.

Lastly, we summarize our findings in Chapter 6 and provide an outlook to possible directions for further research.

Chapter 2.

Background

In this chapter, we briefly discuss the mathematical background of the models and approaches presented in this work. In Section 2.1, we introduce dynamical systems in general and particularly ODEs and PDEs. We further review asymptotic behavior of dynamical systems and methods for stability analysis. Subsequently, we introduce the basic concepts for parameter estimation and uncertainty analysis in Section 2.2. Finally, we summarize ODE and PDE constrained optimization (Section 2.3).

2.1. Dynamical systems

In this thesis, we consider mathematical models for biological processes. For the general theory on dynamical systems we mainly follow (Michel et al., 2015) and (Temam, 1997) adapting their notations for our purposes.

Dynamical systems describe the temporal evolution of a processes by mapping a time t and an initial value u_0 to a state $u(t)$ in some state space V . We will consider deterministic dynamical systems, i.e., systems that yield identical outcomes for identical initial conditions. While there are dynamical systems with discrete time, in this work we focus on systems with continuous time, $t \in I$ for some finite or infinite interval $I \subseteq \mathbb{R}$. The state space, V , is generally a Banach or Hilbert space. It might be finite dimensional as in the case of ODEs or infinite dimensional as in the case of PDEs. Loosely following the notation and definitions in (Michel et al., 2015), we consider temporal evolution based on motions $\Phi(\cdot, u_0, t_0): [t_0, T[\rightarrow V$ with end time point $T > t_0$, $[t_0, T[\in I$ and $\Phi(t_0, u_0, t_0) = u_0$ that map an initial condition u_0 at initial time t_0 to the state of the system at time $t > t_0$. The particular temporal evolution of a dynamical system is described by a family of motions, i.e., a subset of all possible motions. For example, the solutions to an ODE or PDE for

different initial conditions define a family of motions. While based on this definition, motions do not have to be unique, we consider in this work dynamical systems that map an initial state and time to a unique motion.

In the following, we review theoretical aspects of ODEs, PDEs and the asymptotical behavior of dynamical systems for $t \rightarrow \infty$.

2.1.1. Ordinary differential equations

ODE models describe dynamics in a finite dimensional state space and are generally given by

$$\frac{dx}{dt} = f(\theta, t, x), \quad x(t_0) = x_0(\theta), \quad (2.1)$$

with state $x(t) \in \mathbb{R}^{n_x}$, time $t \in [t_0, T] \subset \mathbb{R}$, $T > t_0$, parameter $\theta \in \mathbb{R}^{n_\theta}$, and initial condition $x_0(\theta) \in \mathbb{R}^{n_x}$, where n_x denotes the dimension of the state space and n_θ denotes the dimension of the parameter space. The temporal evolution is defined by the right hand side $f: \mathbb{R}^{n_\theta} \times \mathbb{R} \times \mathbb{R}^{n_x} \rightarrow \mathbb{R}^{n_x}$. The theorem of Picard-Lindelöf guarantees existence and uniqueness of local solutions provided the function f fulfills a Lipschitz condition for $x \in U \subset \mathbb{R}^{n_x}$ for some domain U and is continuous in t (see e.g. Walter (1993)). To ensure local existence, it is sufficient that f is continuous. In terms of dynamical systems, the family of motions generated by the ODE is given by the solution $\Phi(t, x_0, t_0) = x(t, x_0, t_0)$ for all admissible pairs $(t_0, x_0) \in I \times \mathbb{R}^{n_x}$.

A variety of biological processes are modeled using ODEs including biochemical reaction networks, gene regulation, population dynamics and dynamics of infectious diseases (see, e.g., Ingalls (2013); Klipp et al. (2005); Müller and Kuttler (2015)). In most biological applications the system is well behaved and a unique solution exists for all times.

In biological studies, experimental conditions are often regulated by inputs to the system. In contrast to parameters, inputs are usually known and administered to the system from outside, e.g., known concentrations of catalyst in a chemical reaction. Sometimes these inputs, $w(t)$, are explicitly included in the formulation of the system,

$$\frac{dx}{dt} = f(\theta, t, x, w), \quad x(t_0) = x_0(\theta, w), \quad (2.2)$$

to emphasize a dependence of the ODE on the inputs $w(t)$.

2.1.2. Partial differential equations

In addition to ODEs, also PDEs are widely used to model biological processes. Possible uses for PDE models in biology include the description of biochemical processes, e.g., diffusion of molecules in a tissue (Hock et al., 2013), as well as modeling processes in more abstract state spaces like aging in age structured population models (Metzger et al., 2012; Sinko and Streifer, 1967). In this section, we briefly introduce theory for PDE models following (Hinze et al., 2009) and (Zeidler, 1990).

In this work, we consider parabolic PDEs of the form

$$\begin{aligned} u_t &= C(\theta, t, u) + f(\theta, t) \quad \text{on } \Omega_T, \\ u(t_0, x) &= u_0(\theta, x) \quad \text{on } \Omega, \end{aligned} \tag{2.3}$$

where u now denotes the state variable and $x \in \Omega$ denotes the space variable. Further, Ω is an open and bounded set in \mathbb{R}^{n_x} usually with a Lipschitz-boundary and we set $\Omega_T =]t_0, T[\times \Omega$ for the time interval $]t_0, T[\subset \mathbb{R}$. The parameters are finite dimensional and denoted by $\theta \in \mathbb{R}^{n_\theta}$ and $u(t) \in V$ denotes the state in some suitable space V . The dynamics are described by $C: \mathbb{R}^{n_\theta} \times]t_0, T[\times V \rightarrow V^*$ and $f \in V^*$ where V^* is the dual space of V , i.e., the set of linear operators from V to \mathbb{R} (Hinze et al., 2009). The initial condition at $t = t_0$ is given by u_0 . Similarly to ODE models, the solution to (2.3) defines a dynamical system via $\Phi(u, u_0, t_0) = u(t)$ provided a unique solution exists for u_0 .

A concrete class of PDEs that is often used for the description of biological processes are advection-diffusion-reaction equations. The operator C then often takes the form,

$$C(\theta, t, u) = \sum_{i,j=1}^{n_x} (a_{ij}(t, x, \theta)u_{x_i})_{x_j} - \sum_{i=1}^{n_x} b_i(t, x, \theta)u_{x_i} + d(u, \theta), \tag{2.4}$$

where the change of chemical species due to spatial fluxes is described by the advection term, $\sum_{i=1}^{n_x} b_i(t, x, \theta)u_{x_i}$, and the diffusion term, $\sum_{i,j=1}^{n_x} (a_{ij}(t, x, \theta)u_{x_i})_{x_j}$, while the change due to biochemical reactions or degradation is captured in $d(u, \theta)$.

Whereas a parabolic PDE defines a time evolutionary process, stationary processes, i.e., $u_t = 0$, are described by elliptic PDEs,

$$0 = C(\theta, u) + f(\theta) \quad \text{on } \Omega, \tag{2.5}$$

with state $u \in V$, finite dimensional parameters $\theta \in \mathbb{R}^{n_\theta}$ and dynamics described by $C(\theta, u): \mathbb{R}^{n_\theta} \times V \rightarrow V^*$ and $f \in V^*$.

Chapter 2. Background

An example for $C(\theta, u)$ is a diffusion-reaction equation,

$$C(\theta, u) = \sum_{i,j=1}^{n_x} (a_{ij}(x, \theta) u_{x_i})_{x_j} + d(u, \theta), \quad (2.6)$$

with diffusion term $\sum_{i,j=1}^{n_x} (a_{ij}(x, \theta) u_{x_i})_{x_j}$ and where $d(u, \theta)$ describes the change due to biochemical reactions or degradation.

To model a particular biological process using parabolic or elliptic PDEs, the behavior on the domain boundary $\partial\Omega$ has to be specified. We distinguish three common boundary conditions (see e.g. Marin and Öchsner (2019)). Dirichlet boundary conditions give the value at the boundary,

$$u(t, x) = g(x) \quad \text{on } \partial\Omega. \quad (2.7)$$

If the flux through the boundary is prescribed, we obtain von Neumann boundary conditions,

$$\frac{\partial u(t, x)}{\partial \nu} = g(x) \quad \text{on } \partial\Omega. \quad (2.8)$$

Robin boundary conditions describe a mix of flux and value conditions on the boundary,

$$\frac{\partial u(t, x)}{\partial \nu} + \alpha u(t, x) = g(x) \quad \text{on } \partial\Omega. \quad (2.9)$$

Here, $g(x)$ is commonly a continuous function on $\partial\Omega$. Boundary conditions are usually incorporated in the space V , e.g., for Dirichlet conditions with $g = 0$ the space is $V = H_0^1(\Omega)$, or into the domain of C (Amann, 1995).

In contrast to ODEs, there is no unifying theory for existence and uniqueness of solutions for PDEs. For the general parabolic PDE given in (2.3), theory of monotonic operators provides conditions for existence and uniqueness of weak solutions (Zeidler, 1990). For this, we consider separable Hilbert spaces V and H , where V is continuously and densely embedded in H , i.e., $V \subseteq H = H^* \subseteq V^*$ form a Gelfand triple (Hinze et al., 2009), where H is identified with its dual. A common example for a Gelfand triple are the spaces $V = H^1(\Omega)$ and $H = L^2(\Omega)$. In this setting, the initial value is in $u_0 \in H$ and we consider $u \in W(t_0, T) = L^2(t_0, T; V) \cap L^2(t_0, T; V^*)$ (Hinze et al., 2009). Following (Zeidler, 1990, p.770 ff.) existence and uniqueness of a weak solution for fixed θ holds under the following assumptions on the differential operator C :

Assumption 2.1.1.

- The operator $-C(\theta, t, \cdot)$ is monotone and hemicontinuous for each $t \in]t_0, T[$.

2.1. Dynamical systems

- The operator $-C(\theta, t, \cdot)$ is coercive for each $t \in]t_0, T[$, i.e., there exist constants $c_1 > 0$ and $c_2 \geq 0$ such that

$$\langle C(\theta, t, v), v \rangle_{V^*, V} \geq c_1 \|v\|_V^2 - c_2 \quad \text{for all } v \in V, t \in]t_0, T[.$$

- The operator $C(\theta, \cdot, \cdot)$ satisfies a growth condition, i.e., there exist a non-negative function $c_3 \in L^2(t_0, T)$ and a constant $c_4 > 0$ such that

$$\|C(\theta, t, v)\|_{V^*} \leq c_3(t) + c_4 \|v\|_V \quad \text{for all } v \in V, t \in]t_0, T[.$$

- The function $t \mapsto \langle C(\theta, \cdot, w), v \rangle_{V^*, V}$ is measurable on $]t_0, T[$ for all $v, w \in V$.

For the elliptic PDE (2.5), the first two items in Assumption 2.1.1 are sufficient to yield existence and uniqueness of solutions (Zeidler, 1990, p.557 ff.).

2.1.3. Asymptotic behavior

In this section, we consider the asymptotic behavior of general dynamical systems as $t \rightarrow \infty$, using the notation introduced in the beginning of Section 2.1 where $\Phi(\cdot, u_0, t_0)$ now denotes an element of the family of unique motions that correspond to the dynamical system. The state space V is assumed to be a Banach space with norm $\|\cdot\|$. We outline the general definitions of trajectories, ω -limit sets and steady states. We then discuss stability properties as well as methods for stability analysis. This section mainly follows (Michel et al., 2015) and (Temam, 1997).

2.1.3.1. Steady states and stability properties

For a start point $u_0 \in V$ at $t_0 \in I$, we define the trajectory as $\bigcup_{t \geq t_0} \Phi(t, u_0, t_0)$, and the ω -limit set, i.e., the set of points the dynamical system approaches an infinite number of times, as $\omega(u_0) = \bigcap_{s \geq t_0} \overline{\bigcup_{t \geq s} \Phi(t, u_0, t_0)}$. A steady state, $u_e \in V$, also called equilibrium or fixed point, is defined by

$$\Phi(t, u_e, t_0) = u_e \quad \forall t > t_0, t_0 \in I. \quad (2.10)$$

Chapter 2. Background

For example, for ODEs and PDEs steady states correspond to points of vanishing time derivatives,

$$\frac{dx}{dt} = 0 \quad \text{and} \quad u_t = 0. \quad (2.11)$$

The stability properties of a steady state can be classified by the behavior of trajectories in the vicinity.

Definition 2.1.1. (see e.g., (Michel et al., 2015, pp.83 f.))

- a) A fixed point u_e is called a stable steady state if for every $\varepsilon > 0$ and $t_0 \in I$ there exists a $\delta = \delta(\varepsilon, t_0) > 0$ such that from $\|u_0 - u_e\| < \delta$ it follows that $\|\Phi(t, u_0, t_0) - u_e\| < \varepsilon$ for all $t \geq t_0$.
- b) A steady state u_e is asymptotically stable if it is stable and we can find $\delta = \delta(t_0) > 0$ with $\lim_{t \rightarrow \infty} \|\Phi(t, u_0, t_0) - u_e\| = 0$ whenever $\|u_0 - u_e\| < \delta$.
- c) A steady state u_e is uniformly stable if δ in a) is independent of t_0 .
- d) A steady state u_e is exponentially stable if there exists $\gamma > 0$ and for every $\varepsilon > 0$ and every $t_0 \in I$, there exists a $\delta = \delta(\varepsilon) > 0$ such that

$$\|\Phi(t, u_0, t_0) - u_e\| < \varepsilon \exp(-\gamma(t - t_0)) \quad \forall t \geq t_0 \quad (2.12)$$

follows whenever $\|u_0 - u_e\| < \delta$.

- e) If an equilibrium point u_e is not stable, it is called unstable.

We further define attractors and the region of attraction for a dynamical system, see e.g., (Temam, 1997, p.21).

Definition 2.1.2. An attractor is a set $\mathcal{A} \subset V$ that is invariant (i.e., $\Phi(t, u_0, t_0) \in \mathcal{A}$ for all $u_0 \in \mathcal{A}$, $t_0 \in I$ and $t \geq t_0$) and that possesses an open neighborhood U with

$$\text{dist}(\Phi(t, u_0, t_0), \mathcal{A}) \rightarrow 0 \quad \text{as } t \rightarrow \infty \quad (2.13)$$

for every $u_0 \in U$. The neighborhood U might depend on t_0 . The basin of attraction at t_0 is the largest open set U that fulfills this property.

Here, $\text{dist}(\Phi(t, u_0, t_0), \mathcal{A})$ is defined as the infimum of the distance of $\Phi(t, u_0, t_0)$ to the points in \mathcal{A} .

2.1.3.2. Lyapunov functions and stability

Lyapunov functions are a useful tool to analyze stability for both ODEs and PDEs. The theory presented here is based on (Michel et al., 2015, Chapter 9), however we limit the explanations to time independent and differentiable Lyapunov functions.

We consider differential equations in the Banach space V ,

$$\dot{u} = F(u), \quad u(t_0) = u_0, \quad (2.14)$$

where $F: D \subset V \rightarrow V$ and $(t_0, x_0) \in I \times D$. Without loss of generality we assume that the steady state is at the origin, $u_e = 0$. Further let $\dot{\mathcal{V}}(u) = \limsup_{h \rightarrow 0^+} \frac{1}{h} (\mathcal{V}(u + hF(u)) - \mathcal{V}(u))$. Then Theorem 9.2.1 in (Michel et al., 2015) provides the following theory:

Theorem 2.1.1. *If a function $\mathcal{V} \in C(U, \mathbb{R})$ exists for some open set $U \subset D$ containing the fixed point $u_e = 0$ that has the following properties,*

$$(i) \quad \mathcal{V}(u_e) = 0 \text{ and } \mathcal{V}(u) > 0 \text{ for all } u \neq u_e,$$

$$(ii) \quad \dot{\mathcal{V}}(u) \leq 0, \quad \forall u \in U \setminus \{u_e\},$$

u_e is stable.

If (ii) holds with a strict inequality, $\dot{\mathcal{V}}(u) < 0, \forall u \in U \setminus \{u_e\}$, then $u_e = 0$ is uniformly asymptotically stable.

If, additionally, there exist $c_1, c_2, c_3, b > 0$, with $c_1 \|u\|^b \leq \mathcal{V}(u) \leq c_2 \|u\|^b$ and $\dot{\mathcal{V}}(u) \leq -c_3 \|u\|^b$ for all $(u) \in U$, then $u_e = 0$ is exponentially stable.

The proof of the theorem can be found in (Michel et al., 2015).

2.1.3.3. Singular perturbation theory

Dynamical systems can include dynamics on different time scales, e.g., biological processes like diffusion are often much faster than signaling (Hasenauer et al., 2015). However, it might still be worthwhile to include both processes in a model. For systems with different time scales, singular perturbation theory provides a useful method to investigate stability of steady states. As we only employ singular perturbation in the case of ODE models, we limit the discussion to ODEs and follow (Khalil, 1996).

Chapter 2. Background

In singular perturbation theory, we consider the system

$$\begin{aligned}\frac{dx}{dt} &= F(t, x, z, \epsilon) \\ \epsilon \frac{dz}{dt} &= G(t, x, z, \epsilon)\end{aligned}\tag{2.15}$$

for some small $\epsilon > 0$. The states $x \in D_1 \subset \mathbb{R}^{n_x}$ and $z \in D_2 \subset \mathbb{R}^{n_z}$ and differentiable functions $F: \mathbb{R} \times D_1 \times D_2 \times [0, \epsilon_0] \rightarrow \mathbb{R}$ and $G: \mathbb{R} \times D_1 \times D_2 \times [0, \epsilon_0] \rightarrow \mathbb{R}$.

For $\epsilon = 0$, we find the slow system,

$$\frac{dx}{dt} = F(t, x, h(t, x), 0),\tag{2.16}$$

assuming there is a solution $z = h(t, x)$ to $0 = G(t, x, z, 0)$. Using this solution, we can define $\zeta = z - h(t, x)$. Additionally, we introduce the fast time $\tau = \frac{t}{\epsilon}$ and rewrite (2.15) to obtain the fast system

$$\begin{aligned}\frac{dx}{d\tau} &= \epsilon F(t, x, \zeta + h(t, x), \epsilon) \\ \frac{d\zeta}{d\tau} &= G(t, x, \zeta + h(t, x), \epsilon) - \epsilon \frac{\partial h}{\partial t} - \epsilon \frac{\partial h}{\partial x} F(t, x, \zeta + h(t, x), \epsilon).\end{aligned}\tag{2.17}$$

For $\epsilon = 0$, this yields $\frac{dx}{d\tau} = 0$ and the boundary layer system

$$\frac{d\zeta}{d\tau} = G(t, x, \zeta + h(t, x), 0).\tag{2.18}$$

Singular perturbation theory states that under the conditions given in (Khalil, 1996, Theorem 9.3), in essence,

- an isolated root $h(t, x)$ of $0 = G(t, x, z, 0)$ exists,
- the slow system is exponentially stable,
- the boundary layer system is exponentially stable uniformly in (t, x) and
- boundedness of F , G , h and their partial derivatives up to order 2

and for $\epsilon > 0$ sufficiently small the steady state of system (2.15) is exponentially stable. An intuitive interpretation is that the system initially follows close to the fast system moving quickly towards the steady state of the boundary layer system for the given initial values and then moves closely along the slow manifold defined by $0 = G(t, x, z, \epsilon)$ towards the steady state of the slow system (Müller and Kuttler, 2015).

2.2. Statistical models and uncertainty analysis

In the previous section, we introduced ODEs and PDEs as models for biological processes. These models are usually parameter dependent. For the model to give insights into the underlying mechanism or to be useful for predictions and for model selection, the model parameters have to be inferred using experimental data. In this section, we outline how to infer parameters using maximum likelihood estimation and further discuss uncertainty analysis using profile likelihoods. Lastly, we give a brief overview of regularization, its connection to Bayesian inference and model selection.

2.2.1. Maximum likelihood estimation

To infer the model parameters $\theta \in \mathbb{R}^{n_\theta}$ from data, we firstly point out that experiments often do not measure the whole state of the system but only parts. Hence, observation functions or observation operators are used to model the observation process and map the model outputs to the observables. For an ODE the observation function $h: \mathbb{R}^{n_\theta} \times \mathbb{R} \times \mathbb{R}^{n_x} \rightarrow \mathbb{R}^{n_y}$ maps the state x and parameters θ to the observables $y \in \mathbb{R}^{n_y}$

$$y(\theta, t) = h(\theta, t, x(\theta, t)). \quad (2.19)$$

For a parabolic PDE with state u and parameters θ , a mapping $B: \mathbb{R}^{n_\theta} \times]t_0, T[\times V \rightarrow Z$ is used to map to the observables y in some observation space Z , where Z is a Hilbert space,

$$y(\theta, t) = B(\theta, t, u(\theta, t)). \quad (2.20)$$

For an elliptic PDE with state u and parameters θ the mapping $B: \mathbb{R}^{n_\theta} \times V \rightarrow Z$ to $y \in Z$ is used, where the observation space is again a Hilbert space,

$$y(\theta) = B(\theta, u(\theta)). \quad (2.21)$$

The observation function often depends on additional parameters, e.g., scalings or offsets, that have to be likewise estimated together with the dynamic parameters contained in the ODE/PDE.

Regarding the statistical model, we assume that experiments yield noise corrupted measurements of the observables \bar{y} . We then consider time resolved data

$$\mathcal{D} = \left\{ (t_j, \{\bar{y}_{ij}\}_{i=1}^{n_y}) \right\}_{j=1}^{n_t}, \quad (2.22)$$

Chapter 2. Background

where $i \in \{1, \dots, n_y\}$ indicates the component of \bar{y} and $j \in \{1, \dots, n_t\}$ indicates the measurement time point. For a specific noise distribution assumption the parameters are then estimated using a maximum likelihood approach. The likelihood is defined as the probability of observing the data given the parameter vector θ ,

$$\mathcal{L}(\theta) = p(\mathcal{D}|\theta). \quad (2.23)$$

Assuming independent noise, we find

$$\mathcal{L}(\theta) = \prod_{i=1}^{n_y} \prod_{j=1}^{n_t} p(\bar{y}_{ij}|\theta). \quad (2.24)$$

In this work, we focus on two types of measurement noise: additive, normally distributed and multiplicative, log-normally distributed noise. In the case of i.i.d. normally distributed, additive noise, $\bar{y}_{ij} = y_i(t_j) + \epsilon_{ij}$ with $\epsilon_{ij} \sim \mathcal{N}(0, \sigma_{ij}^2)$, the likelihood is given by

$$\mathcal{L}(\theta) = \prod_{i=1}^{n_y} \prod_{j=1}^{n_t} \frac{1}{\sqrt{2\pi}\sigma_{ij}} \exp\left(-\frac{1}{2} \left(\frac{\bar{y}_{ij} - y_i(\theta, t_j)}{\sigma_{ij}}\right)^2\right). \quad (2.25)$$

Considering multiplicative, log-normally distributed noise, i.e., $\bar{y}_{ij} = y_i(t_j)\epsilon_{ij}$ with i.i.d. $\epsilon_{ij} \sim \mathcal{LN}(0, \sigma_{ij}^2)$, yields

$$\mathcal{L}(\theta) = \prod_{i=1}^{n_y} \prod_{j=1}^{n_t} \frac{1}{\sqrt{2\pi}\sigma_{ij}\bar{y}_{ij}} \exp\left(-\frac{1}{2} \left(\frac{\log(\bar{y}_{ij}) - \log(y_i(\theta, t_j))}{\sigma_{ij}}\right)^2\right). \quad (2.26)$$

For this likelihood to map to \mathbb{R} , $\bar{y}_{ij} > 0$ and $y_i(t_j, \theta) > 0$ need to hold. For both noise types, the noise might be parameter-dependent, i.e., $\sigma_{ij}(\theta)$. In this case, the noise parameters have to be estimated together with the other parameters.

The maximum likelihood estimator (MLE) for θ is determined by maximizing the likelihood

$$\theta^* = \arg \max_{\theta \in \Theta} \mathcal{L}(\theta), \quad (2.27)$$

where $\Theta \subset \mathbb{R}^{n_\theta}$ defines the search space. For numerical reasons, usually the equivalent problem of minimizing the negative log-likelihood is considered,

$$\min_{\theta \in \Theta} \{j(\theta) = -\log \mathcal{L}(\theta)\}. \quad (2.28)$$

For additive normally distributed noise this yields

$$j(\theta) = \frac{1}{2} \left(\sum_{i=1}^{n_y} \sum_{j=1}^{n_t} \log(2\pi\sigma_{ij}^2) + \left(\frac{\bar{y}_{ij} - y_i(\theta, t_j)}{\sigma_{ij}} \right)^2 \right), \quad (2.29)$$

and for multiplicative log-normally distributed noise,

$$j(\theta) = \frac{1}{2} \left(\sum_{i=1}^{n_y} \sum_{j=1}^{n_t} \log(2\pi\sigma_{ij}^2 \bar{y}_{ij}^2) + \left(\frac{\log(\bar{y}_{ij}) - \log(y_i(\theta, t_j))}{\sigma_{ij}} \right)^2 \right). \quad (2.30)$$

2.2.2. Uncertainty analysis and confidence intervals

Maximum likelihood estimation yields a point estimator. As measurements of biological processes can be greatly noise-corrupted and usually only part of the system is observed, the MLE is seldom the true parameter. To evaluate the model and model predictions, it is therefore relevant to analyze the uncertainty of parameter estimates and to consider the identifiability of parameters.

The uncertainty is often analyzed by considering confidence intervals. In this work, we use profile likelihood based confidence intervals (Meeker and Escobar, 1995; Raue et al., 2009; Venzon and Moolgavkar, 1988). The profile likelihood for a parameter θ_i is given as

$$PL_i(c) = \max_{\substack{\theta_{j \neq i} \in \Theta \\ \theta_i = c}} \mathcal{L}(\theta), \quad (2.31)$$

i.e., the profile is computed by iteratively fixing parameter θ_i to a certain value c and maximizing the likelihood over the remaining parameters (Murphy and van der Vaart, 2000; Raue et al., 2009).

Using the profile likelihood, we can compute confidence intervals for θ_i ,

$$CI_{\alpha,i} = \left\{ c \mid \frac{PL_i(c)}{\mathcal{L}(\theta^*)} \geq \exp\left(-\frac{\Delta_\alpha}{2}\right) \right\}, \quad (2.32)$$

where Δ_α is the α th-percentile of the χ^2 distribution with one degree of freedom and θ^* is the MLE (Meeker and Escobar, 1995; Raue et al., 2009; Venzon and Moolgavkar, 1988).

Apart from confidence intervals, identifiability of the model parameters is of interest. Parameter identifiability is divided in structural and practical identifiability. Structural identifiability denotes a feature of the model together with the observation function and

is independent of the concrete experimental measurements. It can be analyzed using different approaches (Chis et al., 2011). Practical identifiability analysis on the other hand investigates identifiability of a model including the observation function for a specific incomplete and noisy data set. A parameter is practically identifiable if its confidence interval is bounded. Otherwise, it is non-identifiable (Raue et al., 2009).

2.2.3. Regularization and Bayesian inference

Regularization is often employed in optimization to improve the ill-posedness of optimization problems. Adding a regularization term,

$$j(\theta) = -\log(\mathcal{L}(\theta)) + \rho r(\theta), \quad (2.33)$$

selects for solutions that possess the properties implied by $r(\theta)$. These properties are weighted against the negative log-likelihood using the factor ρ .

This regularization can be interpreted as a prior in a Bayesian sense. In Bayesian analysis not the likelihood but the posterior is considered. The posterior is calculated using Bayes formula

$$p(\theta|\mathcal{D}) = \frac{p(\mathcal{D}|\theta)p(\theta)}{p(\mathcal{D})}, \quad (2.34)$$

for the likelihood $p(\mathcal{D}|\theta) = \mathcal{L}(\theta)$ and a prior $p(\theta)$. As $p(\mathcal{D})$ is constant, it usually suffices to consider $p(\mathcal{D}|\theta)p(\theta)$ for parameter estimation. Similarly to maximum likelihood estimation, the maximum a posterior (MAP) estimator,

$$\theta^{MAP} = \arg \max_{\theta \in \Theta} p(\mathcal{D}|\theta)p(\theta), \quad (2.35)$$

is calculated by minimizing the negative logarithm of the posterior

$$-\log(p(\mathcal{D}|\theta)p(\theta)) = -\log(p(\mathcal{D}|\theta)) - \log(p(\theta)) = -\log(\mathcal{L}(\theta)) - \log(p(\theta)). \quad (2.36)$$

This formulation explains the interpretation of a prior as regularization. Considering for example a multidimensional Gaussian prior $p(\theta) = \mathcal{N}(0, \sigma^2 \text{Id})$ with zero mean and diagonal covariance, we find

$$-\log(p(\mathcal{D}|\theta)p(\theta)) = -\log(\mathcal{L}(\theta)) + \frac{1}{2\sigma^2} \sum_{i=1}^{n_\theta} \theta_i^2 + \text{const}, \quad (2.37)$$

i.e., the Gaussian prior constitutes an l_2 regularization.

2.3. Inverse problems and ODE/PDE constrained optimization

In the Bayesian setting profile likelihoods are replaced by profile posteriors (Hug et al., 2013; Raue et al., 2013a; Vanlier et al., 2012),

$$PP(c) = \max_{\substack{\theta_{j \neq i} \in \Theta \\ \theta_i = c}} p(\theta | \mathcal{D}), \quad (2.38)$$

and these profile posteriors are then applied to quantify uncertainty and compute confidence intervals.

2.2.4. Model selection

Competing hypotheses for the biological mechanism can be evaluated by constructing a mathematical model for each of them. Model selection criteria are then used to determine, how well each model describes the data. If a model can not explain the data sufficiently, the corresponding hypothesis is dismissed. In this work, we consider two model selection criteria, the Akaike information criterion (AIC) (Akaike, 1978) and the Bayesian information criterion (BIC) (Schwarz, 1978).

The AIC is evaluated based on the log-likelihood value at the optimum θ^* and the number of parameters n_θ ,

$$\text{AIC} = -2 \log(\mathcal{L}(\theta^*)) + 2n_\theta. \quad (2.39)$$

The BIC is computed similar to the AIC and includes additionally the number of data points n ,

$$\text{BIC} = -2 \log(\mathcal{L}(\theta^*)) + n_\theta \log n. \quad (2.40)$$

For model selection, the chosen criterion is evaluated for all competing models. The model with the smallest AIC or BIC has the most support. If the difference of the AIC or BIC of a model to the smallest AIC or BIC, respectively, is larger than ten, there is no substantial support for the model according to Burnham and Anderson (2002). If the difference is smaller the model cannot be dismissed.

2.3. Inverse problems and ODE/PDE constrained optimization

In the preceding section, we outlined how to estimate parameters using maximum likelihood estimation. For ODE and PDE models this parameter estimation approach leads to ODE and PDE constrained optimization problems, respectively. In this section, we focus

Chapter 2. Background

on the solution of these optimization problems using global optimization approaches, in particular multi-start local optimization, and gradient based local optimization.

To find the MLE, the negative log-likelihood has to be minimized. In the case of an ODE model this yields an ODE constrained optimization problem,

$$\begin{aligned} \min_{\theta \in \Theta, x \in V} J(\theta, x) \\ \text{s.t. } \dot{x} = f(\theta, t, x), \quad x(t_0) = x_0(\theta). \end{aligned} \quad (2.41)$$

For a parabolic PDE (2.3), we find the optimization problem

$$\begin{aligned} \min_{\theta \in \Theta, u \in V} J(\theta, u) \\ \text{s.t. } \begin{cases} u_t = C(\theta, t, u) + f(\theta, t) & \text{on } \Omega_T, \\ u(t_0, x) = u_0(\theta, x) & \text{on } \Omega, \end{cases} \end{aligned} \quad (2.42)$$

and for an elliptic PDE model (2.5), we find

$$\begin{aligned} \min_{\theta \in \Theta, u \in V} J(\theta, u) \\ \text{s.t. } 0 = C(\theta, u) + f(\theta) \quad \text{on } \Omega. \end{aligned} \quad (2.43)$$

If a unique solution for the ODE/PDE constraints exists for all $\theta \in \Theta$ in the optimization problems (2.41)-(2.43), the optimization problem can be reduced by inserting the solution into the objective function. As an analytical solution to the ODE/PDE might be hard or impossible to compute, a numerical solution can be obtained by simulating the ODE/PDE. For a solution of (2.1), $x(\theta)$, the reduced problem for (2.41) is

$$\min_{\theta \in \Theta} \{j(\theta) = J(\theta, x(\theta))\}. \quad (2.44)$$

Similarly, the reduced problem for (2.42) and (2.43) is

$$\min_{\theta \in \Theta} \{j(\theta) = J(\theta, u(\theta))\} \quad (2.45)$$

where $u(\theta)$ denotes the solution of (2.3) for parabolic constraints and the solution of (2.5) for elliptic constraints. In all cases, the reduced objective function is denoted by $j: \Theta \rightarrow \mathbb{R}$. In general, the afore-mentioned optimization problems are non-convex and non-linear and therefore challenging.

2.3.1. Global optimization

As outlined in the last section, in parameter estimation we are concerned with finding the global minimum of non-convex and non-linear optimization problems. Finding an appropriate optimization algorithm for this class of problems is non-trivial. Possible solutions include, e.g., stochastic algorithms like evolutionary (Bäck, 1996) or particle optimization (Yang, 2010), deterministic algorithms (Raue et al., 2013b) or hybrid optimizers (Balsa-Canto et al., 2008; Vaz and Vicente, 2007). For the type of biological data considered in this work, multi-start local optimization has been shown to perform well (Hross and Hasenauer, 2016; Raue et al., 2013b). Multi-start local optimization performs repeated deterministic local optimizations starting from starting points spread in the search space Θ . The distribution of start points might be uniformly random or based on a latin hypercube scheme. To assess the quality of the multi-start optimization, the results of the individual local optimizations are ordered by the resulting objective function value. If sufficiently many optimizations were performed, different local optima can be found several times and can be identified as plateaus in a plot of the ordered objective function values. Multi-start global optimization then returns the smallest found local minimum as the global minimum. We require a good local algorithm to find the optima consistently on the one hand and to be time efficient on the other hand (Raue et al., 2013b). In order to measure the performance of different local optimizers, we combine these requirements in a metric, the average computation time per converged start, which is computed by dividing the overall computation time for the multi-start optimization by the number of converged starts (Villaverde et al., 2018).

2.3.2. Local optimization

In the previous section, we outlined how to find a global optimum by repeated local optimization. We consider discrete algorithms for local optimization of unconstrained problems based on descent methods (Boyd and Vandenberghe, 2004). For a current (k -th) iterate $\theta^{(k)}$ a descent method chooses the next iterate as

$$\theta^{(k+1)} = \theta^{(k)} + r^{(k)} \Delta\theta^{(k)}, \quad (2.46)$$

with search direction $\Delta\theta^{(k)}$ and step length $r^{(k)}$. For a descent method, we require $j(\theta^{(k+1)}) < j(\theta^{(k)})$. The search direction is often chosen using a steepest descent method,

$$\Delta\theta^{(k)} = \operatorname{argmin}_{\|v\| \leq 1} \nabla j(\theta)^T v \quad (2.47)$$

for some norm $\|\cdot\|$. With the Euclidean norm, we obtain a gradient descent method

$$\Delta\theta^{(k)} = -\nabla j(\theta^{(k)}). \quad (2.48)$$

Provided the Hessian, $H(\theta)$, is positive definite, the corresponding norm can be defined, $\|\theta\|_{H(\theta)}^2 = \theta^T H(\theta)\theta$. Using this norm in (2.47), we find Newton's method (Boyd and Vandenberghe, 2004)

$$\Delta\theta^{(k)} = -H(\theta^{(k)})^{-1}\nabla j(\theta^{(k)}). \quad (2.49)$$

For a given search direction, the step length needs to be chosen. A summary of several methods for this choice can be found in, e.g., (Boyd and Vandenberghe, 2004).

In practice, the parameters are often constraint to be in some search space $\theta \in \Theta$, e.g., using box constraints, $lb_i \leq \theta_i \leq ub_i$ for $i = 1, \dots, n_\theta$ with some bounds $lb_i, ub_i \in \mathbb{R}$. The constrained optimization problem, $\min_{\theta \in \Theta} j(\theta)$, can be solved for example using interior point methods (Byrd et al., 1999, 2000).

2.3.3. Continuous analogues

While discrete iterative local optimization methods as described in the previous section are very prominent, also continuous schemes have been developed (Kose, 1956).

For a given discrete optimization scheme

$$\theta^{(k+1)} = \theta^{(k)} + r^{(k)}g(\theta^{(k)}) \quad (2.50)$$

the continuous analogue is given by

$$\dot{\theta} = g(\theta). \quad (2.51)$$

We have $g(\theta) = -\nabla j(\theta)$ for gradient descent and $g(\theta) = -H^{-1}(\theta)\nabla j(\theta)$ for the Newton method. In Tanabe (1985) it was shown that the continuous analogue tends to exhibit larger regions of attraction and more robust convergence than discrete iterative methods.

2.3.4. Gradients and sensitivities

The local optimization methods listed above all require the computation of the gradient of $j(\theta)$. While the gradient can be approximated using finite differences, it is more efficient

2.3. Inverse problems and ODE/PDE constrained optimization

and accurate to use the analytical derivative (Fröhlich et al., 2017a). Using the chain rule, the analytical derivative in direction of the i -th parameter is given by

$$\frac{dj}{d\theta_i} = \frac{\partial J}{\partial \theta_i} + \frac{\partial J}{\partial x} \frac{\partial x}{\partial \theta_i} \quad (2.52)$$

for problem (2.41) and by

$$\frac{dj}{d\theta_i} = \frac{\partial J}{\partial \theta_i} + \frac{\partial J}{\partial u} \frac{\partial u}{\partial \theta_i}. \quad (2.53)$$

for problems (2.42) and (2.43). The terms $\frac{\partial x}{\partial \theta_i}$ and $\frac{\partial u}{\partial \theta_i}$ are called forward sensitivities for θ_i and describe how the solution of the ODE or PDE changes with the parameter. The sensitivity with respect to parameter θ_i is calculated using the sensitivity equation,

$$\frac{ds_i}{dt} = \frac{\partial}{\partial \theta_i} f(\theta, t, x) + \frac{\partial}{\partial x} f(\theta, t, x) s_i, \quad s_i(t_0) = \frac{\partial}{\partial \theta_i} x_0(\theta) \quad (2.54)$$

for an ODE (2.1),

$$s_{i,t} = \frac{\partial}{\partial \theta_i} C(\theta, t, u) + \frac{\partial}{\partial u} C(\theta, t, u) s_i + \frac{\partial}{\partial \theta_i} f(\theta, t), \quad s_i(t_0) = \frac{\partial}{\partial \theta_i} u_0(\theta) \quad (2.55)$$

for a parabolic PDE (2.3) and

$$0 = \frac{\partial}{\partial \theta_i} C(\theta, u) + \frac{\partial}{\partial u} C(\theta, u) s_i + \frac{\partial}{\partial \theta_i} f(\theta), \quad (2.56)$$

for an elliptic PDE (2.5).

For every parameter, these equations have to be solved together with the original ODE/PDE. Therefore, the size of the ODE/PDE system is expanded by an additional ODE/PDE system of the size of the number of model parameters n_θ increasing the computational complexity.

Chapter 3.

ODE constrained problems with additional steady-state constraints

Many biological processes are naturally regulated to be in a steady state. In this chapter, we develop novel tailored approaches for parameter estimation with single or multiple steady-state constraints for ODE models based on a reduction of the parameter space using the geometry of the steady-state manifold. We first introduce the general biological setting and the corresponding mathematical formulation of the steady-state constrained problem in Section 3.1. Subsequently, we propose two novel estimation approaches in Section 3.2, one based on a retraction operator (Absil et al., 2008) and one based on continuous analogues (Kose, 1956). In Section 3.2.4, we analyze the stability of the continuous analogue. Lastly, in Section 3.3, the new approaches are applied to simulated data and real biological examples to evaluate their properties and compare them with alternative methods.

This chapter is based on and partly identical to the paper *Tailored parameter optimization methods for ordinary differential equation models with steady-state constraints* by Anna Fiedler, Sebastian R ath, Fabian J. Theis, Angelika Hau er and Jan Hasenauer published in BMC Systems Biology, 10(1):80, 2016 (Fiedler et al., 2016), but provides additional results and applications.

3.1. Introduction and mathematical problem formulation

ODE models are often used to describe biological processes, e.g., gene regulation, signal transduction and metabolism (Klipp et al., 2005). In combination with experimental data, these ODE models can give insights into underlying mechanisms. To this end, the model

has to be calibrated to the experimental data, to determine unknown parameters like kinetic rates of the involved reactions or initial conditions (Tarantola, 2005).

Here, we consider parameter estimation for the case that the data is collected from a system that is initially in steady state. Steady states arise naturally in biology (Hass et al., 2019), e.g., processes that are in homeostasis. If the initial state of the system is unknown and has to be inferred from the data, constraining the parameter estimation to pairs of parameters and initial states that fulfill the steady-state condition reduces the search space. These considerations become especially relevant in the context of perturbation experiments.

Due to incomplete observations not all parameters might be inferable from one experiment. Here, perturbation experiments are a valuable tool to acquire informative data. In perturbation experiments cells that are initially in steady state are treated with different levels of external stimuli, e.g., ligands (Bachmann et al., 2011) or small molecules (Bodenmiller et al., 2012). Further external stimuli include knockout experiments (e.g., using CRISPR-Cas9 (Shalem et al., 2015)), or physical stimuli like heat, cold or force. After treatment, the response of the cells is measured (Figure 3.1 A). Depending on the biological process and the input, the stimulus-induced changes might be transient or persistent and either time-course or steady state data of the response are measured. The initial condition of the system corresponds to a stable steady state of the unperturbed system. In the case of persistent changes, the new state is a steady state of the perturbed system. Including the steady-state constraint for the initial condition in the parameter estimation leads to an optimization problem with nonlinear equality constraints.

Even without steady-state constraints, parameter estimation for biological systems is in general a nonlinear and non-convex optimization problem. The additional nonlinear equality constraint increases the complexity. In the following, we will develop two tailored approaches for the parameter estimation in this context.

3.1.1. Formulation of mathematical model

In this section, we describe the general mathematical framework for steady-state constraints in the context of ODE constrained optimization.

In the general case, we consider an ODE

$$\begin{aligned} \frac{dx}{dt} &= f(\theta, t, x), & x(0) &= x_0 \\ 0 &= f_0(\theta, x_0). \end{aligned} \tag{3.1}$$

3.1. Introduction and mathematical problem formulation

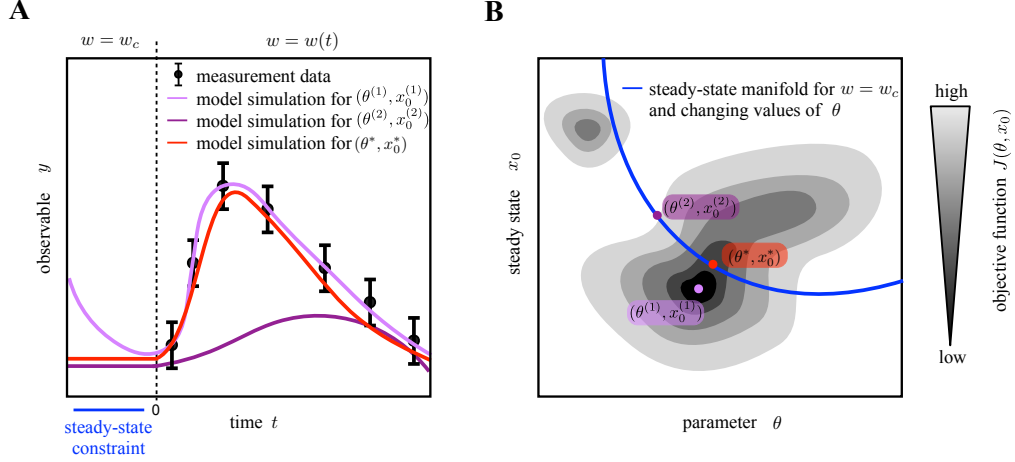


Figure 3.1.: Schematic illustration of optimization problems with steady-state constraint. (A) Measurement data and simulations of the system for three different pairs of parameters and initial conditions: optimum of the unconstrained optimization problem $(\theta^{(1)}, x_0^{(1)})$; suboptimal point on the manifold $(\theta^{(2)}, x_0^{(2)} = \varphi_0(\theta^{(2)}))$; and optimum of the constrained optimization problem $(\theta^*, x_0^* = \varphi_0(\theta^*))$. The system is perturbed at time $t = 0$ and should be in steady state for $t < 0$. (B) Objective function landscape, steady-state manifold and pairs of parameters and initial conditions.

with states $x(t) \in \mathbb{R}^{n_x}$, parameters $\theta \in \mathbb{R}^{n_\theta}$ and initial conditions $x_0 \in \mathbb{R}^{n_x}$ that are given implicitly as the steady state of an ODE with right-hand side $f_0: \mathbb{R}^{n_\theta} \times \mathbb{R}^{n_x} \rightarrow \mathbb{R}^{n_x}$. Hence, the initial condition depends on θ and the form of f_0 . The dynamics of the ODE are given by $f: \mathbb{R}^{n_\theta} \times \mathbb{R} \times \mathbb{R}^{n_x} \rightarrow \mathbb{R}^{n_x}$. In practice, f and f_0 often share some of the parameters θ . To ensure existence and uniqueness of the solution of (3.1), the vector field f is assumed to be Lipschitz continuous. For x_0 to be a meaningful steady state of the biological system described by f_0 , it is assumed to be asymptotically stable in the biologically reasonable parameter regime. The stability of x_0 can be assessed using Lyapunov theory (Khalil, 1996). We denote the collection of all parameter-state pairs (θ, x_0) which fulfill the steady-state constraint as the manifold of steady states. Here, we assume that (3.1) possesses a locally isolated, exponentially stable steady state for parameters θ (cf. Assumptions 3.2.2 and 3.2.3) and that there exists, at least locally, a function $\varphi_0: U_\theta \rightarrow U_x$ with some open subsets $U_\theta \subset \mathbb{R}^{n_\theta}$ and $U_x \subset \mathbb{R}^{n_x}$ which maps the parameters to the corresponding steady state, i.e., $x(0) = \varphi_0(\theta)$ (see Figure 3.1 B). An analytical expression of the function $\varphi_0(\theta)$ is usually not available. As an example, perturbation experiments can be modeled using an initial input w_c and a perturbation input $w(t)$ and considering $f_0(\theta, x) = f_w(\theta, x, w_c)$ and $f(\theta, t, x) = f_w(\theta, x, w(t))$ where f_w describes the dynamics of the studied system.

To describe the measurement process, we denote the observables with $y \in \mathbb{R}^{n_y}$ and the observation function with $h: \mathbb{R}^{n_\theta} \times \mathbb{R} \times \mathbb{R}^{n_x} \rightarrow \mathbb{R}^{n_y}$,

$$y = h(\theta, t, x(\theta, t)). \quad (3.2)$$

3.1.2. Parameter estimation problem

Measurement data is collected for different observables y indexed by $i = 1, \dots, n_y$ and time points t indexed by $j = 1, \dots, N$,

$$\mathcal{D} = \left\{ (t_j, \{\bar{y}_{ij}\}_{i=1}^{n_y}) \right\}_{j=1}^N, \quad (3.3)$$

with number of observables, n_y , and number of measurement time points, N . These measurements are noise-corrupted. For example, additive normally distributed measurement noise, $\epsilon_{ij} \sim \mathcal{N}(0, (\sigma_{ij})^2)$, yields measurements

$$\bar{y}_{ij} = y_i(t_j, \theta, x_0) + \epsilon_{ij}, \quad (3.4)$$

with $y(t, \theta, x_0)$ denoting the solution of (3.1) and (3.2) for parameters θ and initial condition x_0 at time t .

The general optimization problem for parameter estimation from steady state data is of the form

$$\begin{aligned} & \min_{\theta, x_0, x} J(\theta, x_0, x) \\ & \text{s.t.} \quad \begin{cases} y = h(\theta, t, x) \\ \frac{dx}{dt} = f(\theta, t, x), \quad x(0) = x_0 \\ 0 = f_0(\theta, x_0) \end{cases} \end{aligned} \quad (3.5)$$

for some objective function $J: \mathbb{R}^{n_\theta} \times \mathbb{R}^{n_x} \times \mathbb{R}^{n_x} \rightarrow \mathbb{R}$ that measures the goodness-of-fit of model and data. As an example, a maximum likelihood approach for the case of additive normally distributed noise leads to the objective function

$$J(\theta, x_0, x) := \frac{1}{2} \sum_{j=1}^{n_t} \sum_{i=1}^{n_y} \log(2\pi(\sigma_{ij}(\theta))^2) + \left(\frac{\bar{y}_{ij} - y_i(t_j, \theta, x_0)}{\sigma_{ij}(\theta)} \right)^2 \quad (3.6)$$

in which the objective function denotes the negative log-likelihood function, $J(\theta, x_0, x) = -\log p(\mathcal{D}|\theta, x_0, x)$ and the noise parameter, $\sigma_{ij}(\theta)$, is parameter dependent. The solution of (3.5) provides parameter-state pairs $(\hat{\theta}, \hat{x}_0)$ on the steady-states manifold which

3.1. Introduction and mathematical problem formulation

maximize the likelihood. For feasible parameter-state pairs $\hat{x}_0 = \varphi_0(\hat{\theta})$ needs to hold. This formulation can be extended to include multiple experiments and describe parameter estimation for multiple steady states and dose response curves.

3.1.3. Established optimization approaches

Solving optimization problem (3.5) is nontrivial as it is nonlinear and possesses local minima. In the following, we formulate the problem in terms of unconstrained and constrained optimization and describe established approaches for both.

3.1.3.1. Unconstrained optimization

In the unconstrained optimization, an analytical expression of the steady state as a function of the parameter, $\varphi_0(\theta)$, is used to eliminate the constraint and x_0 from optimization problem (3.5), yielding the reduced optimization problem

$$\min_{\theta} \{j(\theta) = J(\theta, \varphi_0(\theta), x(\theta, \varphi_0(\theta)))\} \quad (3.7)$$

which does not possess any equality constraints. E.g., for optimization problem (3.6) the reduced optimization problem is

$$\min_{\theta} \left\{ j(\theta) := \frac{1}{2} \sum_{j=1}^{n_t} \sum_{i=1}^{n_y} \log(2\pi(\sigma_{ij}(\theta))^2) + \left(\frac{\bar{y}_{ij} - y_i(t_j, \varphi_0(\theta), \theta)}{\sigma_{ij}(\theta)} \right)^2 \right\}. \quad (3.8)$$

An available method to compute the steady state is py-substitution (Loriaux et al., 2013; Rosenblatt et al., 2016). In this approach the steady-state equation is not necessarily solved for the state but for the parameters. Using a set of rules based on graph theory, the steady-state equation can be iteratively simplified by choosing the best possible candidate, either a state variable or a parameter, solving the equation for this candidate and removing corresponding equations.

3.1.3.2. Constrained optimization

Alternatively, one can use a constrained optimization method, e.g., an interior point optimization method, to solve the optimization problem (3.5). This is the state-of-the-art method and mostly used in practice.

3.1.4. Problem formulation

Solving optimization problem (3.5) is often challenging, as deterministic local optimization methods have to move along the manifold, resulting in small step-sizes or stagnation, and stochastic local and global optimization methods are only allowed to propose update steps on the manifold, which is not possible in state-of-the-art toolboxes (see, e.g., Egea et al. (2014); Vaz and Vicente (2007)).

Using the reduced formulation and only solving an unconstrained optimization problem is the gold standard solution. However, to use this approach, we need to know the analytical solution to the steady-state constraint, which is not always available.

As we are concerned with the special situation of steady-state constraints, exploiting the structure of the problem might be a potential approach.

Problem 3.1.1. *Given optimization problem (3.5), develop tailored approaches for parameter estimation that exploit the local structure of the problem and investigate the properties of the developed approaches.*

To evaluate the applicability, we want to compare the newly developed tailored approaches to standard constrained and unconstrained local optimization (Section 3.1.3).

3.2. Tailored Algorithms

In this section, we propose tailored algorithms for ODE constrained optimization with steady-state constraints (3.5). The algorithms use the parameter vector θ and the state vector x_0 as optimization variables. First, we consider the first order geometry of the steady-state manifold to develop an update scheme for the state variable for a given parameter update and then present two tailored algorithms for the solution of problem (3.5).

3.2.1. Manifold of steady states

The steady-state constraint defines the steady-state manifold which can be expressed in terms of the mapping $\varphi_0(\theta)$. In the considered setting, the existence of the mapping $\varphi_0(\theta)$ is ensured but an analytical expression is in general not available. For individual parameters θ , the steady state can however be computed by

- solving a feasibility problem (find $x_0 \in \mathbb{R}^{n_x}$ subject to $0 = f_0(\theta, x_0)$),

- simulating the dynamical system until the steady state is reached, or
- combining the simulation of the dynamical system with fine-tuning using the Newton-Raphson method (Shiraishi et al., 2014).

The last two methods are robust and computationally tractable. The computation of the steady state for individual parameters is however not sufficient, as the derivative is also required. To develop a tailored method for solving optimization problem (3.5), we will exploit the first-order geometry of the manifold of steady states. To this end, we consider the sensitivities of the states φ_0 with respect to the parameters θ ,

$$S = \left(s_1, s_2, \dots, s_{n_\theta} \right) \in \mathbb{R}^{n_x \times n_\theta} \quad \text{with} \quad s_i := \frac{\partial \varphi_0}{\partial \theta_i} = \left(\frac{\partial \varphi_{0,1}}{\partial \theta_i}, \dots, \frac{\partial \varphi_{0,n_x}}{\partial \theta_i} \right)^T. \quad (3.9)$$

The dynamics of S are generally governed by the forward sensitivity equation

$$\dot{S} = \frac{\partial f_0}{\partial x_0} S + \frac{\partial f_0}{\partial \theta}. \quad (3.10)$$

In steady state, $\dot{S} = 0$, the equation simplifies to

$$S = - \left(\frac{\partial f_0}{\partial x_0} \right)^{-1} \frac{\partial f_0}{\partial \theta}, \quad (3.11)$$

evaluated at $(\theta, \varphi_0(\theta))$. The invertibility of the Jacobian $(\partial f_0 / \partial x_0)$ follows from local exponential stability of the steady state or can be obtained using the implicit function theorem.

The sensitivity of the steady state with respect to the parameters, S , provides a first-order approximation to $\varphi_0(\theta)$,

$$\varphi_0(\theta + r\Delta\theta) = \varphi_0(\theta) + \nabla_\theta \varphi_0(\theta) r\Delta\theta + o(r). \quad (3.12)$$

The perturbation direction and the step size are denoted by $\Delta\theta$ and r , respectively. Further, o denotes the Landau symbol, defined by: $f: \mathbb{R}^n \rightarrow \mathbb{R}$ is $o(g)$ as $x \rightarrow x_0$ for some $g: \mathbb{R}^n \rightarrow \mathbb{R}$, $g(x) \neq 0$, if $\lim_{x \rightarrow x_0} \frac{f(x)}{g(x)} = 0$. We consider $\hat{\varphi}_0(r) = \varphi_0(\theta + r\Delta\theta)$ and compute the derivative with respect to r , yielding

$$\frac{d\hat{\varphi}_0}{dr} = \nabla_\theta \varphi_0(\theta) \Delta\theta + o(1) = S(\theta + r\Delta\theta, \hat{\varphi}_0(r)) \Delta\theta + o(1) \quad \text{as } r \rightarrow 0. \quad (3.13)$$

This motivates the following update rule for a state variable x_0 on the manifold

$$\frac{dx_0}{dr} = S(\theta, x_0)\Delta\theta. \quad (3.14)$$

Given an update direction $\Delta\theta$ and a step size r , (3.14) provides an update for the steady state corresponding to the parameter $\theta + r\Delta\theta$. Hence, (3.14) enables moves on the steady-state manifold, similar to results in (Absil et al., 2008) for other manifolds.

3.2.2. Hybrid optimization method

In this section, we address Problem 3.1.1 by developing an approach based on the numerical computation of the steady state corresponding to a parameter update (Section 3.2.1) rather than the analytical expression. This is computationally more demanding than using an analytical expression, but also yields the reduced optimization problem (3.7).

This straightforward approach is visualized in Figure 3.2 A. As it can be employed in any state-of-the-art optimization method, we denote it as a hybrid optimization method. Starting at a point (θ^l, x_0^l) , we employ a three-step procedure:

- Step 1: The local optimizer proposes new parameters θ^{l+1} . This yields the point (θ^{l+1}, x_0^l) which is usually not on the manifold of steady states. (Represented as a solid arrow in Figure 3.2 A.)
- Step 2: We compute the steady state x_0^{l+1} for the parameters θ^{l+1} using one of the methods discussed in Section 3.2.1 (with starting points x_0^l). This yields the point $(\theta^{l+1}, x_0^{l+1})$ on the steady-state manifold. (Represented as a dotted arrow in Figure 3.2 A.)
- Step 3: The objective function value $J^{l+1} = J(\theta^{l+1}, x_0^{l+1})$ is computed for parameters θ^{l+1} and numerically calculated steady state x_0^{l+1} . This objective function is provided to the local optimizer. (Not represented in Figure 3.2 A.)

The retraction to the manifold of steady states (Step 2) reduces the problem dimension and eliminates the constraint. To improve the performance of the local optimizer we provide the objective function gradient for the reduced problem (3.7),

$$\frac{dj}{d\theta} = \frac{\partial J}{\partial \theta} + \frac{\partial J}{\partial x_0} \frac{\partial \varphi_0}{\partial \theta} \quad (3.15)$$

with the sensitivity, $(\partial\varphi_0/\partial\theta) = S(\theta, x_0)$, as defined in (3.11).

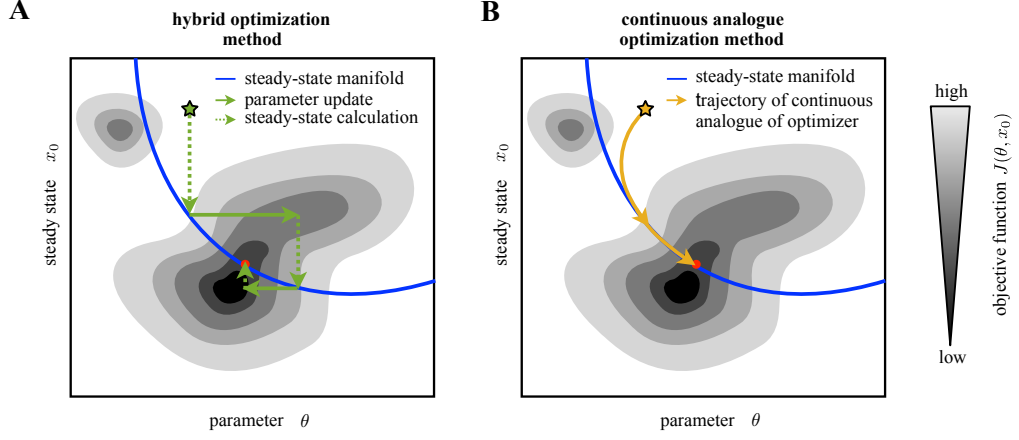


Figure 3.2.: Schematic illustration of (A) the hybrid optimization method and (B) the continuous analogue optimization methods. The path of the optimizers are illustrated along with the manifold of steady states and the objective function contour.

The proposed hybrid optimization method possesses all properties and options of the employed local optimizer. In addition, the retraction accuracy ϵ_{tol} of the convergence criteria $\|f(\theta^l, x_0^l)\|_2 < \epsilon_{\text{tol}}$ has to be selected.

3.2.3. Continuous analogues for optimization

Instead of using a discrete update as in local optimization, one can alternatively consider continuous analogues (illustrated in Figure 3.2 B). The continuous analogue of a gradient descent method is $d\theta/dr = -(dj/d\theta|_{(\theta, x_0)})^T$ (Tanabe, 1985). This ODE system can be coupled with the dynamical system for the state variable x_0 evolving on the steady-state manifold, using $\Delta\theta = -(dj/d\theta|_{(\theta, x_0)})^T$ in (3.14). More generally, we can consider any descent direction $g(\theta, x_0)$ in which j is decreasing. We obtain the ODE system

$$\begin{aligned} \frac{d\theta}{dr} &= g(\theta, x_0) \\ \frac{dx_0}{dr} &= S(\theta, x_0) \frac{d\theta}{dr}, \end{aligned} \quad (3.16)$$

with the steady-state sensitivity $S(\theta, x_0)$. Initialization of (3.16) in a point $(\theta_0, x_{0,0})$ fulfilling (3.1) yields a trajectory evolving on the steady-state manifold, along which the objective function decreases. Formulation (3.16) avoids the repeated simulation-based retractions used by the hybrid optimization method, however it also bears two disadvantages: (i) An appropriate initial point $(\theta_0, x_{0,0})$ has to be determined by solving (3.1); and (ii) numerical errors can result in a divergence from the steady-state manifold. To

address these problems, we introduce the term $\lambda f_0(\theta, x_0)$ which locally retracts the state of the system to the manifold by exploiting the stability properties of the steady state. This yields the system

$$\begin{aligned}\frac{d\theta}{dr} &= g(\theta, x_0) \\ \frac{dx_0}{dr} &= \hat{S}(\theta, x_0) \frac{d\theta}{dr} + \lambda f_0(\theta, x_0).\end{aligned}\tag{3.17}$$

As the retraction term locally draws the trajectory to the steady-state manifold, we no longer require that the initial point $(\theta_0, x_{0,0})$ fulfills the steady-state constraint (3.1). Hence, the Jacobian $\nabla_{x_0} f_0(\theta, x_0)$ might not be invertible. To address this, we define

$$\hat{S}(\theta, x_0) = - \left(\frac{\partial f_0}{\partial x_0} \right)^+ \frac{\partial f_0}{\partial \theta}\tag{3.18}$$

in which $(\partial f_0 / \partial x_0)^+$ denotes the Moore–Penrose pseudoinverse of $(\partial f_0 / \partial x_0)$ at (θ, x_0) . On the steady-state manifold, the Jacobian is invertible and we recover the standard steady-state sensitivity. For a large retraction factor $\lambda \gg 0$, the state (θ, x_0) is retracted quickly to the steady-state manifold, as illustrated in Section 3.3.1.

We consider two possible choices for the descent direction:

- Gradient descent: $g(\theta, x_0) = -\frac{dj}{d\theta}(\theta, x_0)^T = - \left(\frac{\partial J}{\partial \theta}(\theta, x_0) + \frac{\partial J}{\partial x_0}(\theta, x_0) \hat{S}(\theta, x_0) \right)^T$ and
- Newton-type descent: $g(\theta, x_0) = -(\text{FIM}(\theta, x_0) + \mu I)^{-1} \frac{dj}{d\theta}(\theta, x_0)^T$.

In both formulations $\varphi_0(\theta)$ is substituted by x_0 extending the descent direction also to states x_0 not on the steady-state manifold. The Newton-type descent exploits the Fisher Information Matrix (FIM) (Faller et al., 2003) of the reduced optimization problem (3.7). The FIM is a positive semi-definite approximation to the Hessian of the objective function computed from first-order sensitivities. For additive normal distributed noise we find

$$\begin{aligned}\text{FIM}_{kl}(\theta, x_0) &:= \sum_{j=1}^{n_t} \sum_{i=1}^{n_y} -\frac{1}{(\sigma_{ij}^2(\theta))^2} \left(\frac{1}{2} - \frac{(\bar{y}_{ij} - y_i(t_j, \theta, x_0))^2}{\sigma_{ij}^2(\theta)} \right) \frac{\partial \sigma_{ij}^2}{\partial \theta_k}(\theta) \frac{\partial \sigma_{ij}^2}{\partial \theta_l}(\theta) \\ &\quad + \frac{(\bar{y}_{ij} - y_i(t_j, \theta, x_0))}{(\sigma_{ij}^2(\theta))^2} \\ &\quad \left(\frac{\partial y_i}{\partial \theta_k}(t_j, \theta, x_0) \frac{\partial \sigma_{ij}^2}{\partial \theta_l}(\theta) + \frac{\partial \sigma_{ij}^2}{\partial \theta_k}(\theta) \frac{\partial y_i}{\partial \theta_l}(t_j, \theta, x_0) \right) \\ &\quad + \frac{1}{2\sigma_{ij}^2(\theta)} \left(1 - \frac{(\bar{y}_{ij} - y_i(t_j, \theta, x_0))^2}{\sigma_{ij}^2(\theta)} \right) \frac{\partial^2 \sigma_{ij}^2}{\partial \theta_k \partial \theta_l}(\theta) \\ &\quad + \frac{1}{\sigma_{ij}^2(\theta)} \frac{\partial y_i}{\partial \theta_k}(t_j, \theta, x_0) \frac{\partial y_i}{\partial \theta_l}(t_j, \theta, x_0).\end{aligned}\tag{3.19}$$

For parameter independent noise, the expression simplifies to

$$\text{FIM}_{kl}(\theta, x_0) := \sum_{j=1}^{n_t} \sum_{i=1}^{n_y} \frac{1}{\sigma_{ij}^2} \frac{\partial y_i}{\partial \theta_k}(t_j, \theta, x_0) \frac{\partial y_i}{\partial \theta_l}(t_j, \theta, x_0). \quad (3.20)$$

For local optimization of (3.5), the dynamical system (3.17) has to be simulated for $r \rightarrow \infty$. For this, implicit methods with adaptive step-size selection and error control should be employed as (3.17) might be stiff. Appropriate numerical methods are implemented among others in MATLAB and the SUNDIALS package (Hindmarsh et al., 2005). These simulations are stopped as soon as the convergence criterion $\max\{\|d\theta/dr\|, \|dx_0/dr\|\} < \epsilon_{\text{tol}}$ is met.

3.2.4. Asymptotic stability of the optimal parameter-state pair

In the following, we assess the local stability and convergence of the continuous analogue with gradient and with Newton-type descent to a local optimum $(\theta^*, \varphi_0(\theta^*))$ for a single steady-state constraint. We consider the following four conditions.

Assumption 3.2.1. *The minimizer θ^* is a root of the descent direction evaluated on the steady-state manifold, $g(\theta, \varphi_0(\theta))$, i.e., $g(\theta^*, \varphi_0(\theta^*)) = 0$.*

Assumption 3.2.2. *There exists a locally isolated root $\varphi_0(\theta)$ of the equation $f_0(\theta, \varphi_0(\theta)) = 0$ for all $(\theta - \theta^*) \in B_c = \{\tilde{\theta} \in \mathbb{R}^{n_\theta} \mid \|\tilde{\theta}\| \leq c\} \subseteq \mathbb{R}^{n_\theta}$.*

This root is a steady state of the ODE model.

Assumption 3.2.3. *The steady state $\varphi_0(\theta)$ is locally exponentially stable uniformly in θ , with $(\theta - \theta^*) \in B_c$.*

Accordingly, $\exists \zeta_0, \gamma, k > 0$ such that $\|x(t) - \varphi_0(\theta)\| \leq k\|x(0) - \varphi_0(\theta)\| \exp(-\gamma t)$, $\forall (x(0) - \varphi_0(\theta)) \in B_{\zeta_0} := \{\tilde{\Delta}_0(0) \in \mathbb{R}^{n_x} \mid \|\tilde{\Delta}_0(0)\| < \zeta_0\}$ and $\forall (\theta - \theta^*) \in B_c$. This implies that $\exists \zeta \geq \zeta_0$ such that $(x(t) - \varphi_0(\theta)) \in B_\zeta := \{\tilde{\Delta}_0(t) \in \mathbb{R}^{n_x} \mid \|\tilde{\Delta}_0(t)\| < \zeta\}$ (Khalil, 1996).

Assumption 3.2.4. *There exists a neighborhood around θ^* , $(\theta - \theta^*) \in B_c$, in which the descent direction evaluated on the steady-state manifold $\varphi_0(\theta)$, $g(\theta, \varphi_0(\theta))$, is locally uniformly monotonically decreasing.*

Hence, there exists a $\gamma_g > 0$ such that

$$(g(\theta, \varphi_0(\theta)) - g(\theta^*, \varphi_0(\theta^*)))^T (\theta - \theta^*) \leq -\gamma_g \|\theta - \theta^*\|^2 \quad (3.21)$$

for all θ with $(\theta - \theta^*) \in B_c$.

Assumption 3.2.5. *The functions $g(\theta, x_0)$, $f_0(\theta, x_0)$ and $\varphi_0(\theta)$ and their partial derivatives up to order 2 are bounded for $(\theta - \theta^*) \in B_c$ and $(x_0 - \varphi_0(\theta)) \in B_\zeta$.*

Using these assumptions we find:

Theorem 3.2.1. *Let Assumptions 3.2.2 - 3.2.5 be satisfied. Then there exists a λ^* such that for all $\lambda > \lambda^*$ a local minimum (θ^*, x_0^*) of the optimization problem is a locally exponentially stable steady state of the continuous analogue optimization system.*

Proof. To prove Theorem 3.2.1 we use perturbation theory (Khalil, 1996). We define $\varepsilon = \lambda^{-1}$ and shift the optimum to the origin using the linear state transformation,

$$\tilde{\theta} = \theta - \theta^* \quad \text{and} \quad \tilde{x}_0 = x_0 - \varphi_0(\theta^*).$$

This yields the singular perturbed system

$$\begin{aligned} \frac{d\tilde{\theta}}{dr} &= g =: \tilde{F}(\tilde{\theta}, \tilde{x}_0, \varepsilon) \\ \varepsilon \frac{d\tilde{x}_0}{dr} &= \varepsilon \hat{S}g + f_0 =: \tilde{G}(\tilde{\theta}, \tilde{x}_0, \varepsilon), \end{aligned} \tag{3.22}$$

with g , f_0 and \hat{S} evaluated at $(\tilde{\theta} + \theta^*, \tilde{x}_0 + \varphi_0(\theta^*))$. The system (3.22) captures the dynamics of the deviances from the optimal parameter, $\tilde{\theta}$, and the deviances from the steady state for the optimal parameter, \tilde{x}_0 . Furthermore, it possesses the following properties:

- (i) $\tilde{F}(0, 0, \varepsilon) = 0$ and $\tilde{G}(0, 0, \varepsilon) = 0$ as the descent direction vanishes, $0 = g(\theta^*, \varphi_0(\theta^*))$ at the optimum for gradient descent (Assumption 3.2.1) and as the steady-state condition is fulfilled, $0 = f_0(\theta^*, \varphi_0(\theta^*))$, which follows from Assumption 3.2.2.
- (ii) The equation $0 = \tilde{G}(\tilde{\theta}, \tilde{x}_0, 0)$ has the isolated root $\tilde{\varphi}_0(\tilde{\theta}) = \varphi_0(\tilde{\theta} + \theta^*) - \varphi_0(\theta^*)$ (Assumption 3.2.2) with $\tilde{\varphi}_0(0) = 0$.
- (iii) The functions \tilde{F} , \tilde{G} , and φ_0 and their partial derivatives up to order 2 are bounded for $(\tilde{x}_0 - \varphi_0(\tilde{\theta})) \in B_\zeta$ (Assumption 3.2.5).
- (iv) The origin of the reduced systems

$$\frac{d\tilde{\theta}}{dr} = \tilde{F}(\tilde{\theta}, \tilde{\varphi}_0(\tilde{\theta}), \varepsilon) = g(\tilde{\theta} + \theta^*, \tilde{\varphi}_0(\tilde{\theta}) + \varphi_0(\theta^*))$$

for $\varepsilon = 0$ is locally exponentially stable, with $\tilde{\varphi}_0(\tilde{\theta}) + \varphi_0(\theta^*) = \varphi_0(\tilde{\theta} + \theta^*)$ from (ii).

This follows from (Assumption 3.2.4) by considering the Lyapunov function $V(\tilde{\theta}) = \frac{1}{2}\|\tilde{\theta}\|^2$. With this definition, $V(0) = 0$ and $V(\tilde{\theta}) > 0$ for $\tilde{\theta} \neq 0$. Further,

$$\begin{aligned} \frac{dV(\tilde{\theta})}{dr} &= \frac{d}{dr} \tilde{\theta}^T \tilde{\theta} = \tilde{F}(\tilde{\theta}, \tilde{\varphi}_0(\tilde{\theta}), 0)^T \tilde{\theta} = g(\tilde{\theta} + \theta^*, \tilde{\varphi}_0(\tilde{\theta}) + \varphi_0(\theta^*))^T \tilde{\theta} \\ &= g(\theta, \varphi_0(\theta))^T (\theta - \theta^*) = (g(\theta, \varphi_0(\theta)) - g(\theta^*, \varphi_0(\theta^*)))^T (\theta - \theta^*) \\ &\leq -\gamma_g \|\tilde{\theta}\|^2. \end{aligned} \quad (3.23)$$

Hence, the reduced system is exponentially stable (Khalil, 1996, Corollary 3.4).

- (v) The boundary-layer system is derived from (3.22) using the transformation $\tilde{\Delta}_0 = \tilde{x}_0 - \tilde{\varphi}_0(\tilde{\theta}) (= x_0 - \varphi_0(\theta))$, yielding

$$\begin{aligned} \frac{d\tilde{\theta}}{dr} &= g \\ \varepsilon \frac{d\tilde{\Delta}_0}{dr} &= \varepsilon(\hat{S} - S)g + f_0, \end{aligned}$$

with g , f_0 and \hat{S} evaluated at $(\tilde{\theta} + \theta^*, \tilde{\Delta}_0 + \varphi_0(\tilde{\theta} + \theta^*))$ and S evaluated at $(\tilde{\theta} + \theta^*, \varphi_0(\tilde{\theta} + \theta^*))$. After rescaling of the simulation time, $\rho = r/\varepsilon$, and setting $\varepsilon = 0$, we obtain the boundary-layer system

$$\frac{d\tilde{\Delta}_0}{d\rho} = f_0(\tilde{\theta} + \theta^*, \tilde{\Delta}_0 + \varphi_0(\tilde{\theta} + \theta^*)).$$

The origin of this boundary-layer system is locally exponentially stable, uniformly in $\tilde{\theta}$ as the steady state $\varphi_0(\theta)$ of $\frac{dx}{dt} = f_0(x, \theta)$ is exponentially stable uniformly in $\theta = \tilde{\theta} + \theta^*$ (Assumption 3.2.3).

The properties (i)-(v) are the prerequisites of (Khalil, 1996, Theorem 9.3), establishing the existence of $\varepsilon^* > 0$ such that for all $\varepsilon < \varepsilon^*$ systems of type (3.22) are locally exponentially stable. As stability properties are conserved under the performed transformations, we obtain for $\varepsilon = \lambda^{-1}$ Theorem 4.3.1. \square

Stability and convergence are not affected by the approximation of the steady-state sensitivity via \hat{S} . Besides, the theorem can be extended to several steady-state constraints.

Remark 3.2.1. *If the reduced problem j is strongly convex in a neighborhood B_c of θ^* , Assumption 3.2.4 is fulfilled for gradient descent. Strong convexity of j in a neighborhood of θ^* is equivalent to*

$$j(\theta_1) \geq j(\theta_2) + \nabla j(\theta_2)^T (\theta_1 - \theta_2) + \frac{m}{2} \|\theta_1 - \theta_2\|^2 \quad (3.24)$$

for $m > 0$ and all θ_1, θ_2 in the neighborhood of θ^* (Boyd and Vandenberghe, 2004, p. 459). As this hold also conversely,

$$j(\theta_2) \geq j(\theta_1) + \nabla j(\theta_1)^T(\theta_2 - \theta_1) + \frac{m}{2}\|\theta_1 - \theta_2\|^2, \quad (3.25)$$

we find by adding the two equations and reformulating

$$(\nabla j(\theta_1) - \nabla j(\theta_2))^T(\theta_1 - \theta_2) \geq m\|\theta_1 - \theta_2\|^2. \quad (3.26)$$

With $g(\theta) = -\nabla j(\theta)$, we find

$$(g(\theta_1) - g(\theta_2))^T(\theta_1 - \theta_2) \leq -m\|\theta_1 - \theta_2\|^2. \quad (3.27)$$

As this holds for all $\theta_1, \theta_2 \in B_c$ it also holds for $\theta_2 = \theta^*$.

Remark 3.2.2. For Newton-type descent, $g(\theta) = -(\text{FIM}(\theta) + \mu I)^{-1}\nabla j(\theta)$, Assumption 3.2.4 is equivalent to

$$(-(\text{FIM}(\theta) + \mu I)^{-1}\nabla j(\theta) + (\text{FIM}(\theta^*) + \mu I)^{-1}\nabla j(\theta^*))^T(\theta - \theta^*) \leq -\gamma_g\|\theta - \theta^*\|^2. \quad (3.28)$$

Using $\nabla j(\theta^*) = 0$, we reformulate

$$\begin{aligned} (-\text{FIM}(\theta) + \mu I)^{-1}\nabla j(\theta)^T(\theta - \theta^*) &\leq -\gamma_g\|\theta - \theta^*\|^2 \\ (-\nabla j(\theta))^T(\text{FIM}(\theta) + \mu I)^{-T}(\theta - \theta^*) &\leq -\gamma_g\|\theta - \theta^*\|^2. \end{aligned} \quad (3.29)$$

With symmetry of $(\text{FIM}(\theta) + \mu I)^{-1}$ and $\nabla j(\theta^*) = 0$ we arrive at

$$-(\nabla j(\theta) - \nabla j(\theta^*))^T(\text{FIM}(\theta) + \mu I)^{-1}(\theta - \theta^*) \leq -\gamma_g\|\theta - \theta^*\|^2. \quad (3.30)$$

Using the equivalence of norms in finite dimensions, we can write

$$\langle \nabla j(\theta) - \nabla j(\theta^*), \theta - \theta^* \rangle_{(\text{FIM}(\theta) + \mu I)^{-1}} \geq \gamma_g C(\theta) \|\theta - \theta^*\|_{(\text{FIM}(\theta) + \mu I)^{-1}}^2, \quad (3.31)$$

where the scalar product is induced by $(\text{FIM}(\theta) + \mu I)^{-1}$, $\|\cdot\|_{(\text{FIM}(\theta) + \mu I)^{-1}}$, is the corresponding norm and $C(\theta)$ is the equivalence constant. Hence, j has to fulfill Assumption 3.2.4 using the scalar product induced by $(\text{FIM}(\theta) + \mu I)^{-1}$ and the corresponding norm instead of the standard scalar product for all $\theta \in B_c$.

3.3. Application and evaluation

In the following, we illustrate the behavior of the proposed parameter estimation approaches. Furthermore, the performance of the proposed methods will be compared to standard constrained and unconstrained optimization methods. For this purpose, we consider simulation examples for which the ground truth is known. Furthermore, we test the methods in two application examples using real data.

3.3.1. Conversion process

In this section, we illustrate the proposed optimization methods by studying parameter estimation for a conversion process from steady-state data. Conversion processes are a common motif in biological systems and provide a simple test case.

We considered the conversion of a substance A to a substance B with rate constant θ_1 that can be modified by an input $w \in \mathbb{R}_+$ and the reverse conversion with rate constant θ_2 ,



with parameters $\theta = (\theta_1, \theta_2) \in \mathbb{R}_+^2$. Assuming mass conservation ($[A] + [B] = 1$) and mass action kinetics (Ingalls, 2013), the temporal evolution of the concentration of biochemical species A, $x = [A]$, is governed by

$$\begin{aligned} \frac{dx}{dt} &= \theta_2 - (\theta_1 w + \theta_2)x, & x(0) &= x_0 \\ y &= x \end{aligned} \quad (3.33)$$

with $x_0 \in \mathbb{R}_+$ denoting the initial concentration. The unique steady state of model (3.33) for given input w is

$$\varphi_0(\theta) = \frac{\theta_2}{\theta_1 w + \theta_2}. \quad (3.34)$$

To illustrate the properties of the hybrid and the continuous analogue methods, we considered the estimation of the parameters θ from artificial time-resolved data for y . The artificial data at the time points $t_j = [0, 0.1, 0.5, 1, 2]$ were obtained by simulation of (3.33) for $\theta = (4, 1)$ and $w = 0.4$ starting from the steady state of the system with control condition $w_c = 1$ at $t = 0$. Assuming unit variance for observation errors, yielded the optimization

problem

$$\min_{\theta, x_0} \left\{ J(\theta, x_0) := \frac{1}{2} \sum_{j=1}^N (\bar{y}_j - y(t_j, \theta, x_0))^2 \right\} \quad (3.35)$$

s.t. $0 = \theta_2 - (\theta_1 + \theta_2)x_0,$

in which \bar{y}_j denotes the measured concentration of A at time point t_j and $y(t_j, \theta, x_0)$ denotes the solution of (3.33) for initial conditions $x(0) = x_0$ and input $w = 0.4$ at time point t_j .

Illustration of hybrid optimization method The hybrid optimization method evaluates the steady state numerically but exploits the gradients of the objective functions (3.15).

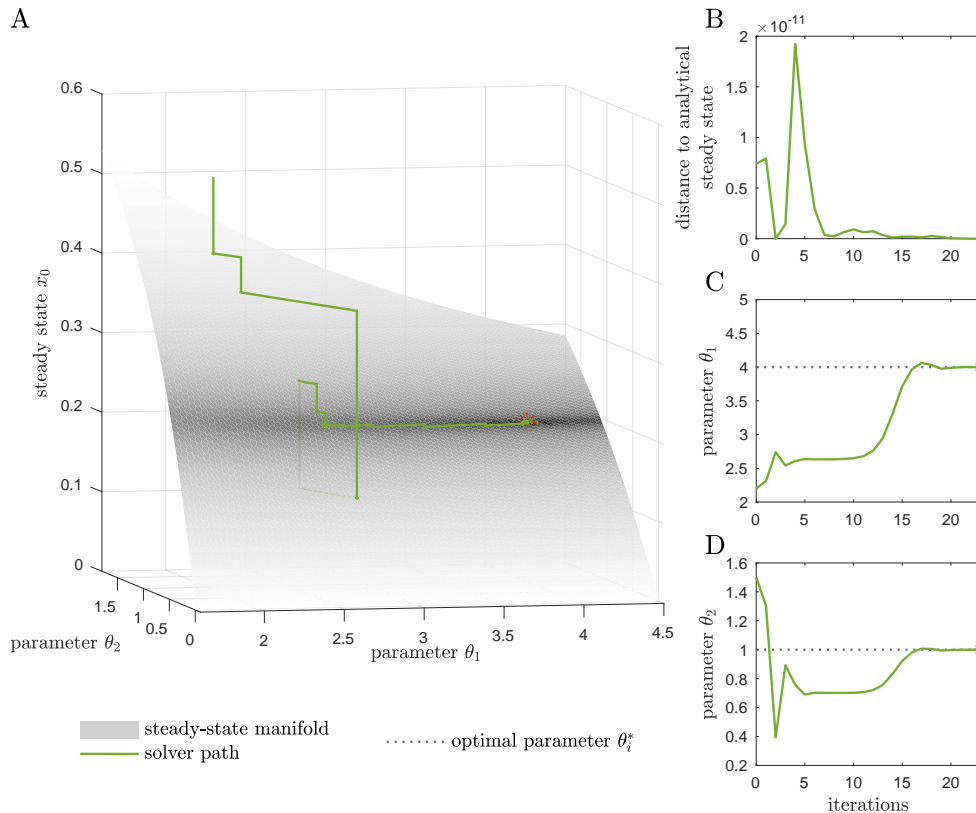


Figure 3.3.: Illustration of the hybrid optimization method for the conversion process (3.33). (A) Path of the hybrid optimization method (full lines), the true optimum (red star) and the steady-state manifold (surface, (3.34)) are shown. The objective function values are indicated by the surface coloring where darker colors represents smaller objective function values. The solver path is partially covered by the steady-state manifold. (B) The path of the steady state and (C), (D) the parameters (full lines), their endpoints (stars), and optimal parameter value respective true steady state (3.34) for the parameters (dotted lines) are depicted.

3.3. Application and evaluation

To this end, the objective function gradient,

$$\frac{dJ}{d\theta} = - \sum_{j=1}^N (\bar{y}_j - y(t_j, \theta, x_0)) \left(\frac{\partial y(t_j, \theta, x_0)}{\partial \theta} + \frac{\partial y(t_j, \theta, x_0)}{\partial x_0} \frac{\partial \varphi_0}{\partial \theta} \right), \quad (3.36)$$

and the local sensitivities of the steady state for $w = 1$,

$$S(\theta, x_0) = \left(\frac{\partial \varphi_0}{\partial \theta_1}, \frac{\partial \varphi_0}{\partial \theta_2} \right) = \left(\frac{-x_0}{\theta_1 + \theta_2}, \frac{1 - x_0}{\theta_1 + \theta_2} \right), \quad (3.37)$$

were derived. The local sensitivities depend merely on derivatives of the vector field and can be computed without knowledge of an analytical expression of the steady state.

The trajectory of the hybrid optimization method is illustrated in Figure 3.3. At the end of each iteration, the simulation-based retraction ensured that the parameter-state pair is on the steady-state manifold (Figure 3.3 A and B). On the steady-state manifold, the optimizer reached a narrow valley for θ_1 and θ_2 and then moved along the valley to the optimum (Figure 3.3 A, C and D). The behavior was similar for other starting points.

Illustration of continuous analogue optimization method For the second method the continuous analogue of the gradient descent method was derived. This yields the dynamical system

$$\begin{aligned} \frac{d\theta_1}{dr} &= - \sum_{j=1}^N (\bar{y}_j - y(t_j, \theta, x_0)) \left(\frac{\partial y(t_j, \theta, x_0)}{\partial \theta_1} + \frac{\partial y(t_j, \theta, x_0)}{\partial x_0} s_1 \right) \\ \frac{d\theta_2}{dr} &= - \sum_{j=1}^N (\bar{y}_j - y(t_j, \theta, x_0)) \left(\frac{\partial y(t_j, \theta, x_0)}{\partial \theta_2} + \frac{\partial y(t_j, \theta, x_0)}{\partial x_0} s_2 \right) \\ \frac{dx_0}{dr} &= \left(\frac{-x_0}{\theta_1 + \theta_2}, \frac{1 - x_0}{\theta_1 + \theta_2} \right) \left(\frac{d\theta_1}{dr}, \frac{d\theta_2}{dr} \right)^T + \lambda(\theta_2 - (\theta_1 u + \theta_2)x_0), \end{aligned} \quad (3.38)$$

with initial conditions $\theta_1(0) = \theta_{1,0}$, $\theta_2(0) = \theta_{2,0}$ and $x_0(0) = x_{0,0}$ and s_1 and s_2 given by (3.37). It can be verified that the objective function J is locally strictly convex in θ – the parameters are locally identifiable – and that the model (3.33) is asymptotically stable. Accordingly, system (3.38) converges to a local optimum of the constrained optimization problem (3.35).

To illustrate the continuous analogue optimization method, we simulated the continuous analogue of the gradient descent method. Exemplary trajectories are depicted in Figure 3.4. We found that for retraction factors $\lambda > 0$, the states $(\theta_1, \theta_2, x_0)^T$ converged to the optimal solution. As retraction renders the steady-state manifold (3.34) attractive,

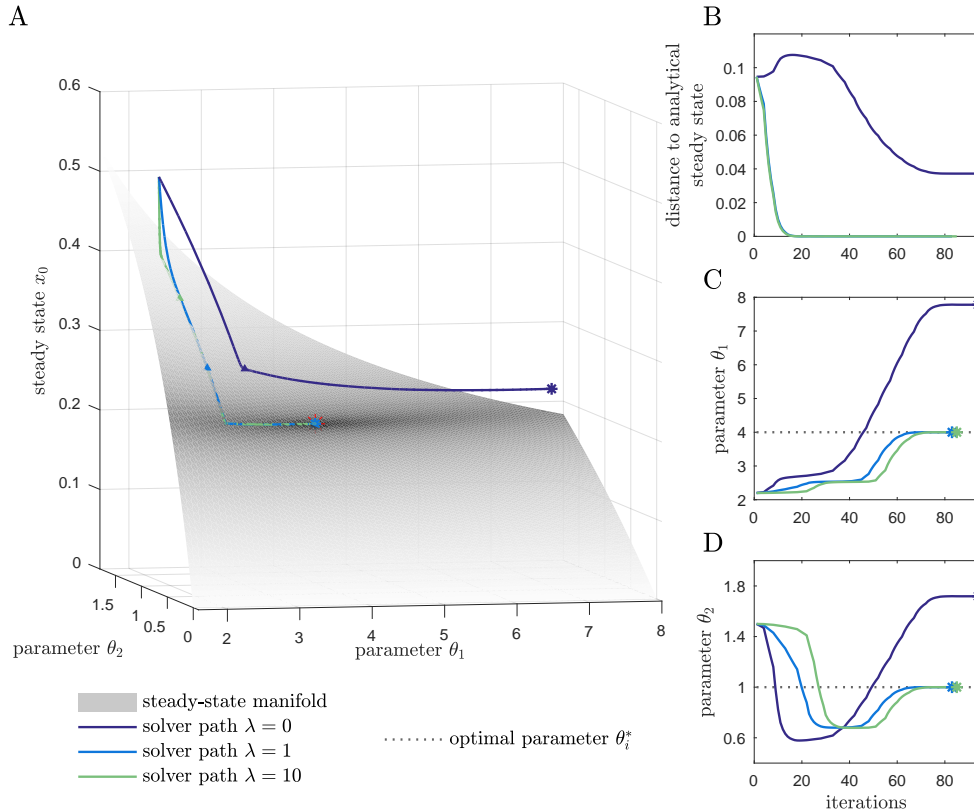


Figure 3.4.: Illustration of the continuous analogue optimization method for the conversion process (3.33). (A) Trajectory of the continuous analogue for different retraction factors λ (full lines), the endpoints (stars), the true optimum (red star), and the steady-state manifold (surface, (3.34)) are shown. The objective function values are indicated by the surface coloring where darker colors represents smaller objective function values. (B) The path of the steady state and (C), (D) the parameters (full lines), their endpoints (stars), and optimal parameter value respective true steady state (3.34) for the parameters (dotted lines) are depicted.

also for initial conditions $(\theta_{1,0}, \theta_{2,0}, x_{0,0})^T$ which do not fulfill the steady-state condition, fast convergence to the steady-state manifold could be achieved using $\lambda \gg 1$ (Figure 3.4 A and D). For large retractions ($\lambda \gg 1$), the dynamic consisted of two phases: (Phase 1) the state x_0 converged quickly to the parameter-dependent steady state (3.34) (Figure 3.4 A and B); and (Phase 2) the state $(\theta_1, \theta_2, x_0)^T$ moved along the steady-state manifold to the global optimum (Figure 3.4 A, C and D).

3.3.2. Behavior in the presence of bistability, bifurcations and oscillations

Biological systems often possess multiple stable steady states (Gardner et al., 2000; Ozbudak et al., 2004) and non-trivial ω -limit sets, e.g., stable limit cycles (Kholodenko, 2000).

To evaluate the continuous analogue optimization method and the hybrid optimization method, we considered a bistable system (with hysteresis) and a system with a Hopf bifurcation.

3.3.2.1. Bistable system

To study the performance of the proposed methods in the presence of bistability we considered the system

$$\dot{x} = \frac{x^{20}}{1+x^{20}} - x + \theta, \quad x(0) = x_0. \quad (3.39)$$

as described in (Müller and Kuttler, 2015, Chapter 5, p. 532). We assumed direct observation of the state and the objective function $J(\theta, x_0) = (\bar{x}_0 - x_0)^2$ with $\bar{x}_0 \approx 1.5$. The measurement \bar{x}_0 is the larger of the two steady states for $\theta = 0.5$. The resulting minimization problem is

$$\begin{aligned} \min_{\theta, x_0} J(\theta, x_0) &= (\bar{x}_0 - x_0)^2 \\ \text{s.t. } 0 &= \frac{x_0^{20}}{1+x_0^{20}} - x_0 + \theta. \end{aligned} \quad (3.40)$$

The optimization was performed using the hybrid and the continuous analogue method. While the system exhibits saddle node bifurcations and bistability, Assumptions 3.2.2 - 3.2.5 are fulfilled locally around the optimal parameter-state pair. However, it is still interesting to observe how the approaches behave for different start points. Figure 3.5 illustrates the convergence and 5 example trajectories. We observed a higher convergence for the hybrid method compared to the continuous analogue, as the continuous analogue diverges for some starting points. This occurs when the retraction term cancels the update in the state direction.

Nevertheless both methods overall achieved good convergence to the optimal point, illustrating good properties in the presence of bistability.

3.3.2.2. System with Hopf bifurcation

To study the performance of the proposed methods for a system that exhibits stable oscillations in the parameter region considered for optimization, we consider the system

$$\dot{x}_1 = -x_2 + x_1(\theta - x_1^2 - x_2^2), \quad x_1(0) = x_{1,0} \quad (3.41)$$

$$\dot{x}_2 = x_1 + x_2(\theta - x_1^2 - x_2^2), \quad x_2(0) = x_{2,0}. \quad (3.42)$$

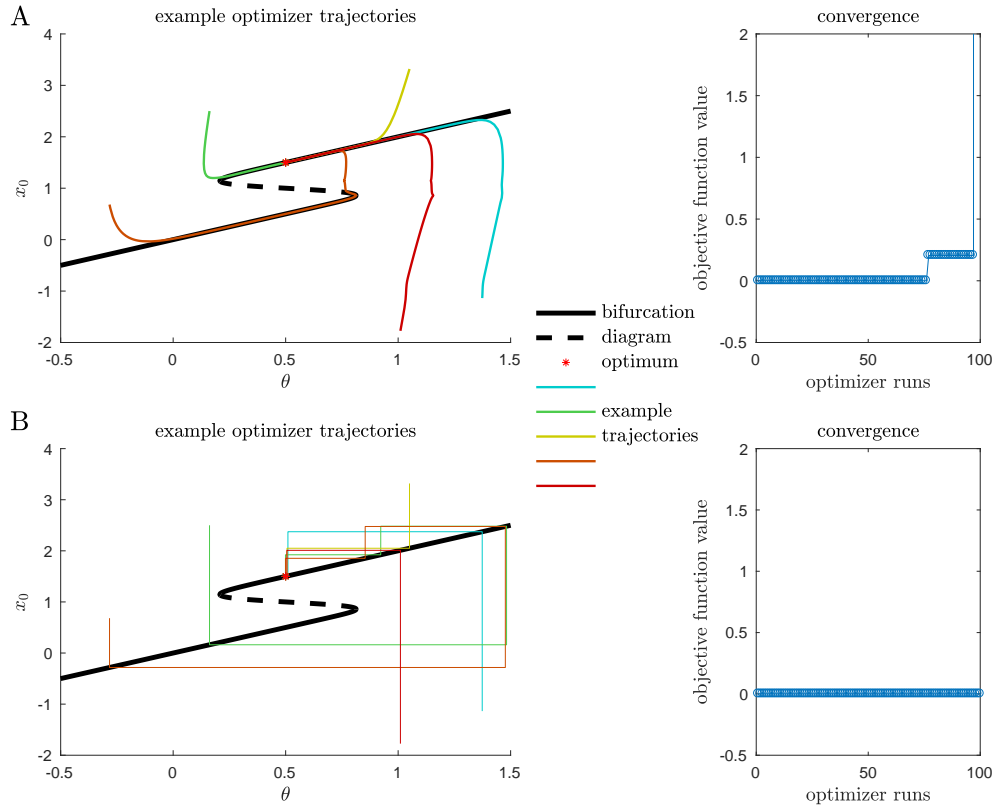


Figure 3.5.: Evaluation of (A) the continuous analogue optimization method and (B) the hybrid optimization method for bistable system. (left) 5 example optimizer trajectories and (right) results of 100 local optimization runs.

This is the normal form for systems with Hopf bifurcation as described in (Müller and Kuttler, 2015, Chapter 2, p. 227). Hopf bifurcations occur in biological applications for example in the context of the cAMP oscillations (Sgro et al., 2015). For the simulated data, the measurement for (3.41) was taken to be $x_0 = (0, 0)^T$ which corresponds to the steady state for $\theta \leq 0$. For $\theta > 0$ this system exhibits stable limit cycles and no stable steady state. We considered a least squares objective function yielding the optimization problem

$$\begin{aligned} \min_{\theta, x_0} J(\theta, x_0) &= \sum_{i=1}^2 x_{0,i}^2 \\ \text{s.t. } 0 &= x_{0,2} + x_{0,1}(\theta - x_{0,1}^2 - x_{0,2}^2) \\ 0 &= x_{0,1} + x_{0,2}(\theta - x_{0,1}^2 - x_{0,2}^2). \end{aligned} \tag{3.43}$$

The optimization was performed using only the continuous analogue optimization method as the simulation in the hybrid optimization step is not converging for $\theta > 0$ as no stable steady state exists. Figure 3.6 illustrates example optimizer trajectories. Even though

3.3. Application and evaluation

Assumption 3.2.4 is violated for this example as $j(\theta)$ is only convex, but not strictly convex, the results still illustrate a good convergence of the continuous analogue optimization method to the set of parameters θ with steady state $x_0 = (0,0)^T$, independent of the starting point as long as the starting point lies outside the limit cycle. For starting points inside the limit cycle, the system diverges. However, convergence strongly depends on the chosen retraction factor λ . If the retraction is very strong, trajectories tend to get trapped close to the stable limit cycle, if started in the basin of attraction of the limit cycle.

In summary, the analyses of these simple systems indicated that both methods also facilitate the analysis of models with multiple stable steady states and that the continuous analogue can also be employed for systems exhibiting limit cycles in the considered parameter regime.

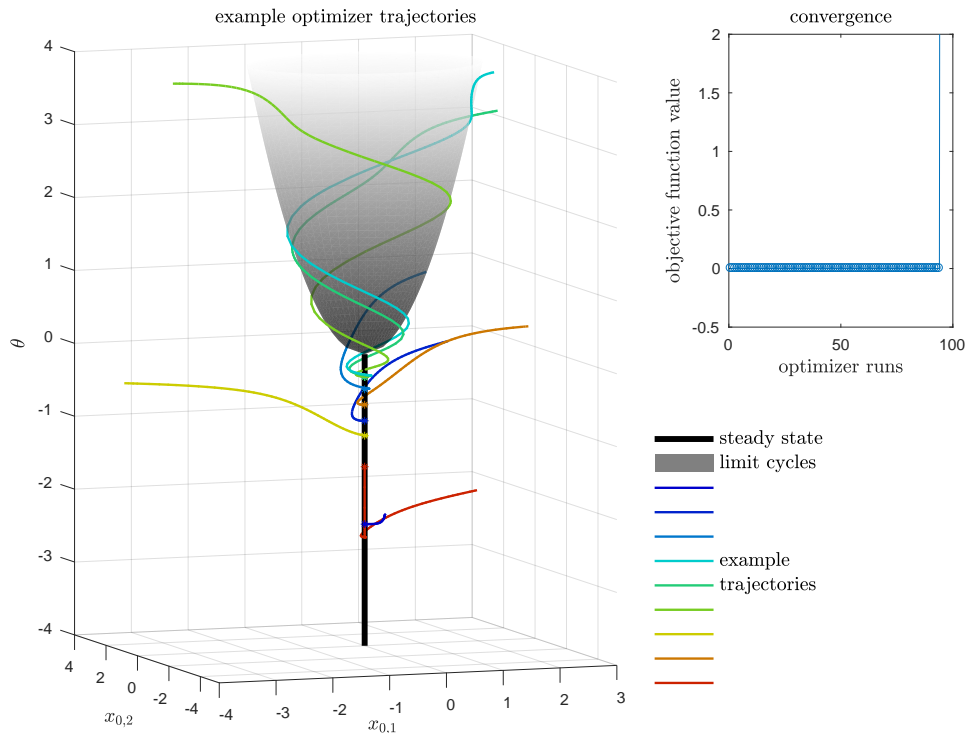


Figure 3.6.: Evaluation of the simulation-based optimization method for the system with Hopf-bifurcation. (left) 10 example optimizer trajectories and (right) results of 100 local optimization runs.

3.3.3. NGF-induced ERK signaling in primary sensory neurons

After applying the proposed methods on simulated data, we now consider experimental data for biological processes. To evaluate and compare established and proposed local optimization methods for problems with steady-state constraints, we analyze nerve growth factor (NGF)-induced extracellular signal-regulated kinases (ERK) phosphorylation in primary sensory neurons. The stimulus NGF binds to cellular receptors and induces the ERK phosphorylation (Andres et al., 2010). This modulates neuronal activity by triggering ion channel phosphorylation and protein expression (Nicol and Vasko, 2007).

Growth-factor induced ERK signaling is a potential target for novel pain therapies (Andres et al., 2013) and therefore of high practical relevance. In addition, this application is well-suited for the evaluation of the proposed methods as NGF dose-response curves at late time points have been recorded where equilibration of the system can be assumed. These data provide multiple steady-state constraints for the thorough assessment of the methods. In the following, we will compare the performance of unconstrained, constrained, hybrid and continuous analogue optimization in the presence of multiple steady-state constraints.

3.3.3.1. Experimental data for NGF-induced ERK phosphorylation

ERK phosphorylation in response to different concentrations of NGF was previously quantified using quantitative automated microscopy (Andres et al., 2010). This technique provides single-cell data from which population average data can be derived. These population average data are highly reproducible and quantitative but provide merely the relative ERK phosphorylation in comparison to the control as no calibration curve is employed. The unknown scaling constant is denoted by s .

3.3.3.2. Mathematical model of NGF-induced ERK phosphorylation

NGF induces ERK phosphorylation by binding to the NGF receptor tropomyosin receptor kinase A (TrkA). The complex TrkA:NGF activates the protein Ras which in turn phosphorylates the protein Raf. Phospho-Raf (pRaf) phosphorylates the protein MEK and phospho-MEK (pMEK) phosphorylates ERK (Figure 3.7). We use the model introduced

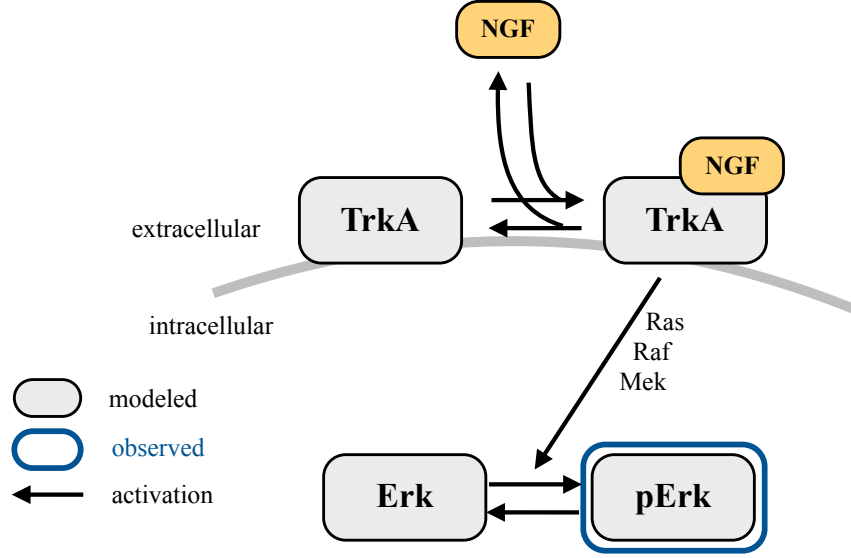



Figure 3.7.: Schematic of the model considered for NGF-induced ERK phosphorylation in primary sensory neurons.

in (Hasenauer et al., 2014), that accounts for five reactions:



These reactions describe binding of NGF to TrkA (R_1 and R_2), TrkA:NGF-mediated ERK phosphorylation (R_3), basal ERK phosphorylation (R_4) and ERK dephosphorylation (R_5). Brackets indicate the concentration of a biochemical species. The input is the initial NGF concentration, $u = [\text{NGF}]_0$, and the measured output is the relative pERK concentration,

$$y = s[\text{pERK}]. \tag{3.45}$$

The experimental noise ϵ is assumed to be normally distributed with the unknown variance σ^2 , $\epsilon \sim \mathcal{N}(0, \sigma^2)$.

In accordance with previous publications, we assume conservation of mass ($[\text{NGF}] + [\text{TrkA:NGF}] = [\text{NGF}]_0$, $[\text{TrkA}] + [\text{TrkA:NGF}] = [\text{TrkA}]_0$ and $[\text{ERK}] + [\text{pERK}] = [\text{ERK}]_0$) and excess of NGF ($[\text{NGF}]_0 \gg [\text{TrkA}]_0$) (see Hasenauer et al. (2014) and references

therein). Under these assumptions the activities of the NGF receptor TrkA and ERK are captured by

$$\begin{aligned}\frac{dx_1}{dt} &= k_1 u (k_3[\text{TrkA}]_0 - x_1) - k_2 x_1, \\ \frac{dx_2}{dt} &= (x_1 + k_4) (s[\text{ERK}]_0 - x_2) - k_5 x_2, \\ y &= x_2,\end{aligned}\tag{3.46}$$

in which $x_1 = k_3[\text{TrkA}:\text{NGF}]$ and $x_2 = s[\text{pERK}]$. This model possesses a minimal number of parameters, $\theta = (k_1, k_2, k_3[\text{TrkA}]_0, k_4, s[\text{ERK}]_0, k_5, \sigma^2)$ and is structurally identifiable.

The parameter- and input-dependent steady state of (3.46) is given by

$$\begin{aligned}\varphi_{0,1}(\theta, w) &= k_3[\text{TrkA}]_0 \frac{k_1[\text{NGF}]_0}{k_1[\text{NGF}]_0 + k_2}, \\ \varphi_{0,2}(\theta, w) &= s[\text{ERK}]_0 \frac{\varphi_{0,1}(\theta, w) + k_4}{\varphi_{0,1}(\theta, w) + k_4 + k_5}.\end{aligned}\tag{3.47}$$

This steady state exists for all positive parameters and is exponentially stable.

3.3.3.3. Parameter estimation problem with multiple steady-state constraints

In this study, the unknown parameters $\theta \in \mathbb{R}_+^7$ and the states $x_{0,1}$ and $x_{0,2}$ for each considered input of NGF were inferred from published dose response data (Andres et al., 2010) using ML estimation. The data set contains six different NGF doses, yielding an optimization problem with $7 + 2 \cdot 6 = 19$ optimization variables and $2 \cdot 6 = 12$ nonlinear equality constraints. This nonlinear optimization problem was solved using multi-start local optimization. The local optimization was performed using unconstrained, constrained and hybrid optimization as well as continuous analogue optimization using gradient and Newton-type descent. In practice, we added the diagonal of $\text{FIM}(\theta, x_0)$ and the maximal diagonal value to the FIM, $\text{diag}(\text{FIM}(\theta, x_0)) + \max_i(\text{FIM}_{ii}(\theta, x_0))$ instead of μI for the Newton-type descent. This proved to be more stable, but was chosen entirely heuristically. Bounds and properties for the parameters and initial conditions are provided in Table 3.1.

To assess the convergence properties, the constraint satisfaction/violation and the computation time, the local optimization methods were initialized with the same 100 sampled starting points. The results are summarized in Figure 3.8. Additionally, we assessed the dependence of the convergence properties on λ (Figure 3.9).

3.3. Application and evaluation

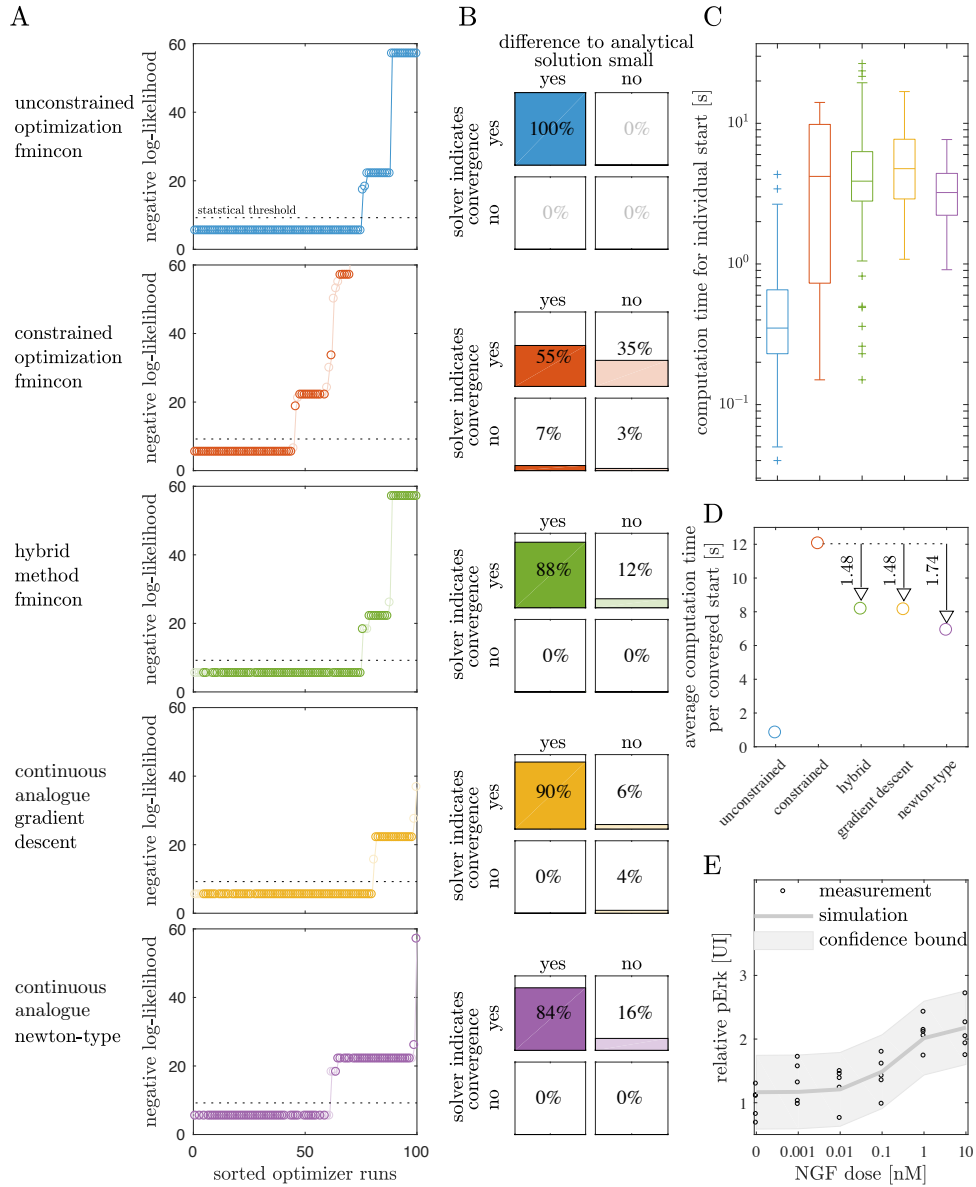


Figure 3.8.: Comparison of optimization methods for the NGF-induced ERK activation model. (A) Final objective function values (darker color in steady-state, lighter color not in steady-state), (B) comparison of convergence criteria with respect to steady-state constraint (i.e., the difference to the analytical solution), (C) computation time for 100 runs and (D) average computation time per converged start of unconstrained optimization method (using the MATLAB optimization method `fmincon`), constrained optimization algorithm (`fmincon`), the hybrid optimization method and the proposed continuous analogue optimization methods with gradient descent and Newton-type updates are depicted. (E) The best fit to the data is shown.

Table 3.1.: Lower bounds, upper bounds and scale used for fitting of the parameters of the model for NGF-induced ERK phosphorylation.

	Parameter name	Lower bound	Upper bound	Scale used for optimization	Bound used for optimization	Bound used for sampling
Kinetic parameters	k_1	10^{-5}	$10^{3.5}$	\log_{10}	yes	yes
	k_2	10^{-5}	$10^{3.5}$	\log_{10}	yes	yes
	$k_3[\text{TrkA}]_0$	10^{-5}	$10^{3.5}$	\log_{10}	yes	yes
	k_4	10^{-5}	$10^{3.5}$	\log_{10}	yes	yes
	$s[\text{ERK}]_0$	10^{-5}	$10^{3.5}$	\log_{10}	yes	yes
	k_5	10^{-5}	$10^{3.5}$	\log_{10}	yes	yes
Initial conditions	$x_{0,1}$	0	3	linear	no	yes
	$x_{0,2}$	0	3	linear	no	yes

3.3.3.4. The convergence properties of unconstrained, hybrid optimization and continuous analogue are comparable

To assess the convergence of the optimization methods, we sorted and visualized the objective function values achieved in the individual optimizer starts (Figure 3.8 A). In addition, we determined the percentage of converged starts. A start is considered to be converged if the final point cannot be rejected compared to the ML estimate using the likelihood ratio test with a significance level of 0.05.

As expected, we found that the gold standard – the unconstrained optimization method – showed the best convergence properties. It converged in 75% of the starts to the global optimum. A similar convergence was achieved by the proposed methods, hybrid optimization and continuous analogue optimization using gradient descent. The third proposed method – continuous analogue optimization using Newton-type descent – displayed intermediate convergence properties (60% of the starts converged to the global optimum). The state-of-the-art method – constrained optimization – exhibited the poorest convergence. It converged in 45% of the starts. Hence, the proposed optimization methods are superior to constrained optimization regarding convergence to the global optimum.

Beyond differences in the convergence to the global optimum, the convergence to local optima differed. The results of unconstrained, constrained and hybrid optimization revealed three local optima. The local optima with the worst objective function values are hardly found using continuous analogue optimization, indicating altered regions of attraction.

3.3.3.5. Hybrid optimization and continuous analogue provide reliable estimates of the steady states

The individual optimization methods enforce the steady-state constraints differently. What all methods have in common is that the steady-state constraint $f(\theta, x_0, w_c) = 0$ is relaxed to a constraint on the norm of the vector field, i.e., $\|f(\theta, x_0, w_c)\|_2 < \epsilon_f$. Accordingly, parameter-state pairs returned by the optimization methods do not fulfill steady-state constraints exactly. Different optimization methods might even achieve different accuracies. In addition, a bound for the difference of the estimated steady state x_0 for a parameter θ and the true steady state $\varphi_0(\theta)$, $\Delta x_0 = x_0 - \varphi_0(\theta)$, is usually not available.

We studied the relation of the solver indicating convergence based on the vector field ($\|f\|_2 < \epsilon_f$) and the difference of the estimated to the analytical steady-state being small ($\|\Delta x_0\|_2 < \epsilon_x$) for the different optimization methods. In our opinion a good optimizer should achieve equivalence of the two criteria. This would mean that enforcing the constraint of the vector field ensures a good approximation of the steady state. The result is depicted in Figure 3.8 B for a tolerance of 10^{-6} for both ϵ_f and ϵ_x .

The unconstrained optimization used an analytical expression of the steady state and therefore the two criteria are identical. Hybrid and continuous analogue optimization also achieved a good agreement of both criteria, with $\sim 85\%$. In $\sim 15\%$ of the cases, the solver indicated convergence based on the vector field constraint but the steady-state estimate is off ($\|\Delta x_0\|_2 > \epsilon_x$). The precise percentage depended heavily on the retraction factor λ for the simulation-based optimization method. For the constrained optimization, all possible combinations were observed and the two criteria agreed in merely 55% of the runs. In summary, the results indicate that the proposed methods provide reliable estimates for the steady states while constrained optimization yields many inconsistent parameter-state pairs.

3.3.3.6. Hybrid optimization and continuous analogue are faster than constrained optimization

Based on the computation times for individual starts and the average computation time per converged start (Figure 3.8 C and D), unconstrained optimization using a (usually not available) analytical expressions for steady states was most efficient. The individual runs were fast and the percentage of converged starts was high. Hybrid and continuous analogue methods were roughly 10 times slower but these methods can be applied if analytical expressions for steady states are not available. Furthermore, these methods

were 1.5 times faster than constrained optimization due to the improved convergence rate. Additionally, the fit to the data for the optimal parameters is convincing (Figure 3.8 E). Accordingly, we conclude that hybrid and continuous analogue optimization are promising approaches in the presence of multiple steady-state constraints.

3.3.3.7. Retraction factor

As an example, we studied the influence of the retraction factor λ on the convergence and computation time for NGF-induced ERK phosphorylation. To this end, we considered retraction factors varying over several orders of magnitude and performed for each of these λ a multi-start optimization with 100 starts using continuous analogue optimization with gradient descent and Newton-type descent directions. The results are illustrated in Figure 3.9. In the case of gradient descent, the computation time decreased with increasing λ , which also resulted in decreasing average computation time per converged start. Choosing $\lambda > 2000$ could not decrease the average computation time further. It could also be observed that there seems to be an upper limit on the maximal computation time. This is caused by a restriction on the maximal possible function evaluations allowed. In the case of a Newton-type descent, increasing λ also decreased the computation time, however, not as drastically as when using gradient descent. Again, also the average computation time per converged start decreased, but plateaued at much smaller values of λ than in the gradient descent case.

3.3.4. Raf/MEK/ERK signaling in HeLa cells after release from S-phase arrest

The analysis of published data in the previous section was complemented by teaming up with the group of Angelika Hausser (University of Stuttgart) to study Raf/MEK/ERK signaling in HeLa cells after release from S-phase arrest. Experimental studies revealed that cell-cycle is, among others, controlled by Raf/MEK/ERK signaling (Chambard et al., 2007; Zhang and Liu, 2002). The signaling dynamics in different cell-cycle phases as well as the cell-cycle-dependent relevance of feedback mechanisms (Fritsche-Guenther et al., 2011) are however still not completely unraveled although a more thorough understanding could provide valuable insights into treatment resistance (Fritsche-Guenther et al., 2011). Using the new data and model selection we studied the relevance of negative feedback from phospho-ERK to Raf activation during G1/S phase transition.

3.3. Application and evaluation

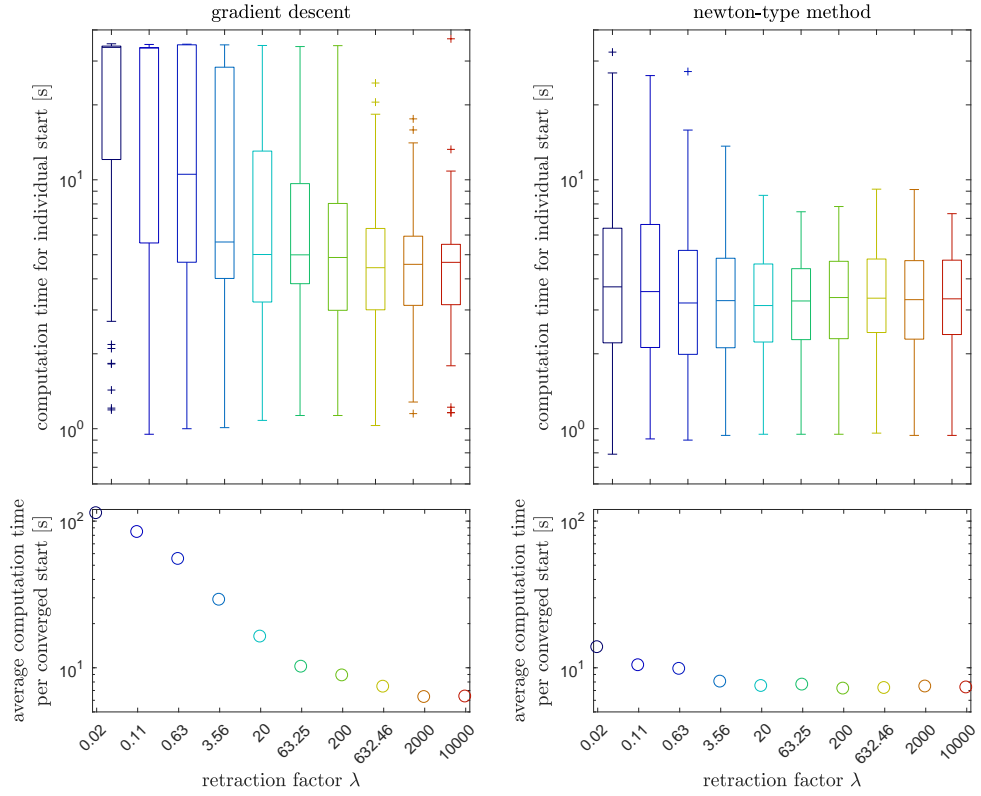


Figure 3.9.: Comparison of the influence of retraction factor λ on the computation time and the average computation time per converged start for the continuous analogue optimization method using gradient descent and Newton-type descent directions for the NGF-induced ERK activation model.

In addition to its biological relevance, the Raf/MEK/ERK pathway is well-suited for the evaluation of the proposed optimization methods and the comparison to state-of-the-art methods. The pathway is nonlinear, yielding a nonlinear and non-convex optimization problem. Furthermore, we consider a synchronized cell population which reached a steady state before the start of the experiment. Accordingly, a steady-state constraint has to be enforced and fitted along with time-resolved data for perturbation experiments.

3.3.4.1. Experimental data for Raf/MEK/ERK signaling after release from S-phase arrest

To study the Raf/MEK/ERK pathway, HeLa cells were synchronized and arrested in cell cycle between cell cycle states G1 and S. The arrest was lifted and the dynamics of phospho-MEK and phospho-ERK were quantified using Western blotting. This experiment was repeated after treatment with the inhibitors Sorafenib and UO126 to explore the dynamic

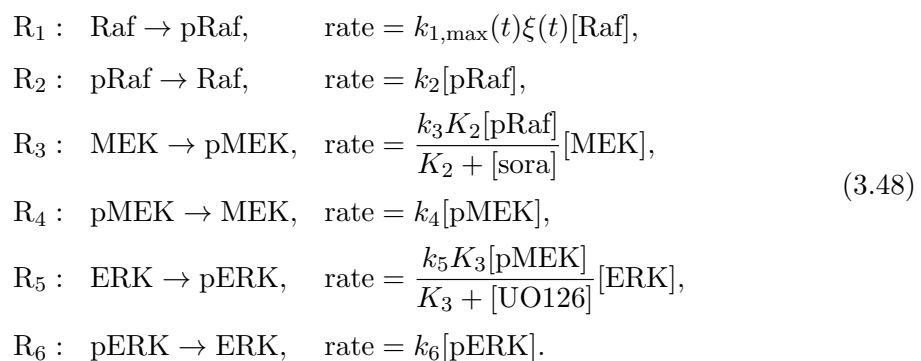
range of the pathway. Sorafenib is an inhibitor of Raf kinases (Wilhelm et al., 2008) and UO126 is a highly selective inhibitor of MEK (Favata et al., 1998).

As Western blots are merely semi-quantitative, they only provide the relative activity of phospho-MEK and phospho-ERK at different time points and under different conditions. The unknown scaling constants differ between blots and measured species. For a detailed discussion of characteristics of Western blot data we refer to (Kreutz et al., 2007).

3.3.4.2. Mathematical model for Raf/MEK/ERK signaling after release from S-phase arrest

Raf/MEK/ERK signaling is induced by myriads of intra- and extracellular signals (Chambard et al., 2007; Kholodenko, 2007). These signals converge on the level of Raf kinase, which they phosphorylate. The phosphorylated Raf kinase phosphorylates MEK, which in turn phosphorylates ERK. Phosphorylated ERK induces downstream signaling and can downregulate the Raf activity (Santos et al., 2007). The latter establishes a negative feedback loop (Fritsche-Guenther et al., 2011; Kholodenko, 2000). The activity of Raf and MEK can be inhibited by Sorafenib and UO126, respectively. The pathway is illustrated in Figure 3.10.

We developed a model for Raf/MEK/ERK signaling which accounts for the core proteins as well as their inhibition with Sorafenib and UO126. The model considers six reactions:



The upstream signaling is summarized in the time-dependent rate constant $k_{1,\max}(t)$ with the flexible parameterization

$$k_{1,\max}(t) = k_{1,0} + k_{1,1} \left(1 - e^{-\frac{t}{\tau_1}}\right) e^{-\frac{t}{\tau_2}} \tag{3.49}$$

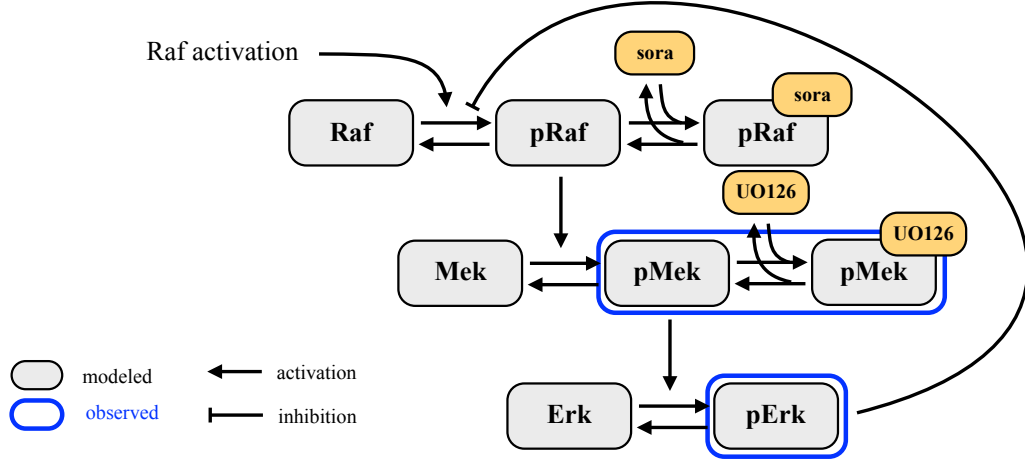


Figure 3.10.: Schematic of the model considered for the Raf/MEK/ERK signaling after release from S-phase arrest.

(as proposed in the Data2Dynamics toolbox (Raue et al., 2015)). The effects of Sorafenib and UO126 are captured by a reduction in the kinase activity of pRaf and pMEK (R_3 and R_5), where K_2 and K_3 are dissociation constants for the binding of the inhibitors (Ingalls, 2013).

Experimental studies showed an inhibition of Raf phosphorylation by pERK (Fritsche-Guenther et al., 2011). This feedback is however context-dependent (Santos et al., 2007). To study the importance of this feedback during the G1/S phase transition, we considered two model hypotheses:

H1: Inhibition of Raf phosphorylation by pERK: $\xi(t) = \frac{K_1}{K_1 + [pERK]}$

H2: No inhibition: $\xi(t) = 1$

We assumed that Raf, MEK and ERK are conserved ($[Raf] + [pRaf] = [Raf]_0$, $[MEK] + [pMEK] = [MEK]_0$ and $[ERK] + [pERK] = [ERK]_0$), yielding the ODE model

$$\begin{aligned}
 \frac{d[pRaf]}{dt} &= k_{1,\max}(t)\xi(t)([Raf]_0 - [pRaf]) - k_2[pRaf] \\
 \frac{d[pMEK]}{dt} &= \frac{k_3 K_2 [pRaf]}{K_2 + [sora]} ([MEK]_0 - [pMEK]) - k_4 [pMEK] \\
 \frac{d[pERK]}{dt} &= \frac{k_5 K_3 [pMEK]}{K_3 + [UO126]} ([ERK]_0 - [pERK]) - k_6 [pERK],
 \end{aligned} \tag{3.50}$$

with experimental input $w = ([\text{sora}], [\text{UO126}])$. The four collected Western blots, $b = 1, \dots, 4$, provide time-resolved relative data for different experimental inputs,

$$\begin{aligned} y_{1,b} &= s_{1,b}[\text{pMEK}] \\ y_{2,b} &= s_{2,b}[\text{pERK}] \end{aligned} \quad (3.51)$$

with unknown, blot-dependent scaling constants $s_{1,b}$ and $s_{2,b}$. The experimental noise for $y_{1,b}$ and $y_{2,b}$ is assumed to be normally distributed and proportional to the scaling constant, $\epsilon_{1,b} \sim \mathcal{N}(0, s_{1,b}^2 \sigma_1^2)$ and $\epsilon_2 \sim \mathcal{N}(0, s_{2,b}^2 \sigma_2^2)$. The unknown variances are denoted by σ_1^2 and σ_2^2 .

For H1 and H2, model (3.50)-(3.51) possesses more than 20 parameters. Several of these parameters are structurally non-identifiable, including the absolute abundances of Raf, MEK and ERK. To circumvent these non-identifiabilities, we reformulate the model in terms of the fractions of phosphorylated proteins: $x_1 = [\text{pRaf}]/[\text{Raf}]_0$, $x_2 = [\text{pMEK}]/[\text{MEK}]_0$ and $x_3 = [\text{pERK}]/[\text{ERK}]_0$. This yields

$$\begin{aligned} \frac{dx_1}{dt} &= k_{1,\max}(t)\xi(t)(1 - x_1) - k_2x_1 \\ \frac{dx_2}{dt} &= \frac{k_3[\text{Raf}]_0K_2x_1}{K_2 + [\text{sora}]}(1 - x_2) - k_4x_2 \\ \frac{dx_3}{dt} &= \frac{k_5[\text{MEK}]_0K_3x_2}{K_3 + [\text{UO126}]}(1 - x_3) - k_6x_3 \\ y_{1,b} &= s_{1,b}[\text{MEK}]_0x_2 \\ y_{2,b} &= s_{2,b}[\text{ERK}]_0x_3 \end{aligned} \quad (3.52)$$

with blot index $b = 1, \dots, 4$ and rescaled experimental noise $\tilde{\epsilon}_{1,b} \sim \mathcal{N}(0, s_{1,b}^2 \sigma_1^2 / [\text{MEK}]_0^2)$ and $\tilde{\epsilon}_{2,b} \sim \mathcal{N}(0, s_{2,b}^2 \sigma_2^2 / [\text{ERK}]_0^2)$. The reformulated model does not depend explicitly on the total abundances $[\text{Raf}]_0$, $[\text{MEK}]_0$ and $[\text{ERK}]_0$ but only on products and ratios of these parameters with other parameters, e.g., $k_3[\text{Raf}]_0$. Defining these products and ratios as new parameters eliminates non-identifiabilities and reduces the number of parameters. As all parameters are non-negative, a log-parameterization is used for parameter estimation (Raue et al., 2013b). The states of the reformulated model are between 0 and 1. In addition to the kinetic, scaling and noise parameters, the initial conditions of the models for H1 and H2 are unknown. However, as the cells are arrested in S-phase with $k_{1,\max}(0) = k_{1,0}$ and $w = 0$, the initial conditions are the corresponding steady states. After significant manual preprocessing of the steady-state constraints, analytical expressions $\varphi_0(\theta)$ for the steady states as a function of the other parameters were calculated with symbolic math toolboxes. The results for models H1 and H2 can be found in the

3.3. Application and evaluation

Table 3.2.: Lower bounds, upper bounds and scale used for fitting of the parameters of model for Raf/MEK/ERK signaling.

	Parameter name	Lower bound	Upper bound	Scale used for optimization	Bound used for optimization	Bound used for sampling
Kinetic parameters	$k_{1,1}/k_{1,0}$	10^{-4}	10^2	log	yes	yes
	τ_1	10^{-4}	10^2	log	yes	yes
	τ_2	10^{-4}	10^2	log	yes	yes
	$K_1/[\text{ERK}]_0$	10^{-4}	10^2	log	yes	yes
	$k_{1,1}$	10^{-4}	10^2	log	yes	yes
	k_2	10^{-4}	10^2	log	yes	yes
	K_2	10^{-4}	10^2	log	yes	yes
	$k_3[\text{Raf}]_0$	10^{-4}	10^2	log	yes	yes
	k_4	10^{-4}	10^2	log	yes	yes
	K_3	10^{-4}	10^2	log	yes	yes
Scaling parameters	$k_5[\text{MEK}]_0$	10^{-4}	10^2	log	yes	yes
	k_6	10^{-4}	10^2	log	yes	yes
Noise parameters	$s_{1,b}[\text{MEK}]_0$	10^{-4}	10^2	log	yes	yes
	$s_{2,b}[\text{ERK}]_0$	10^{-4}	10^2	log	yes	yes
Initial conditions	σ_1^2	10^{-10}	10^2	log	yes	yes
	σ_2^2	10^{-10}	10^2	log	yes	yes
Initial conditions	$x_{0,1}$	0	1	linear	no	yes
	$x_{0,2}$	0	1	linear	no	yes
	$x_{0,3}$	0	1	linear	no	yes

Appendix A. The properties of the parameters for the model of Raf/MEK/ERK signaling after release from S-phase arrest are provided in Table 3.2.

3.3.4.3. Parameter estimation problem with multiple perturbation data sets

We inferred the model parameters and initial conditions from the Western blot data using ML estimation. The data set provides time-resolved data for three conditions (control & two perturbations), all starting from the same steady-state. The optimization problem was solved using multi-start local optimization. The local optimization was performed using unconstrained, constrained and hybrid optimization as well as continuous analogue optimization using gradient and Newton-type descent. Each method started at the same points. As in the previous example, we added the diagonal of $\text{FIM}(\theta)$ and the maximal diagonal value to the FIM, $\text{diag}(\text{FIM}(\theta, x_0)) + \max_i(\text{FIM}_{ii}(\theta, x_0))$ instead of μI .

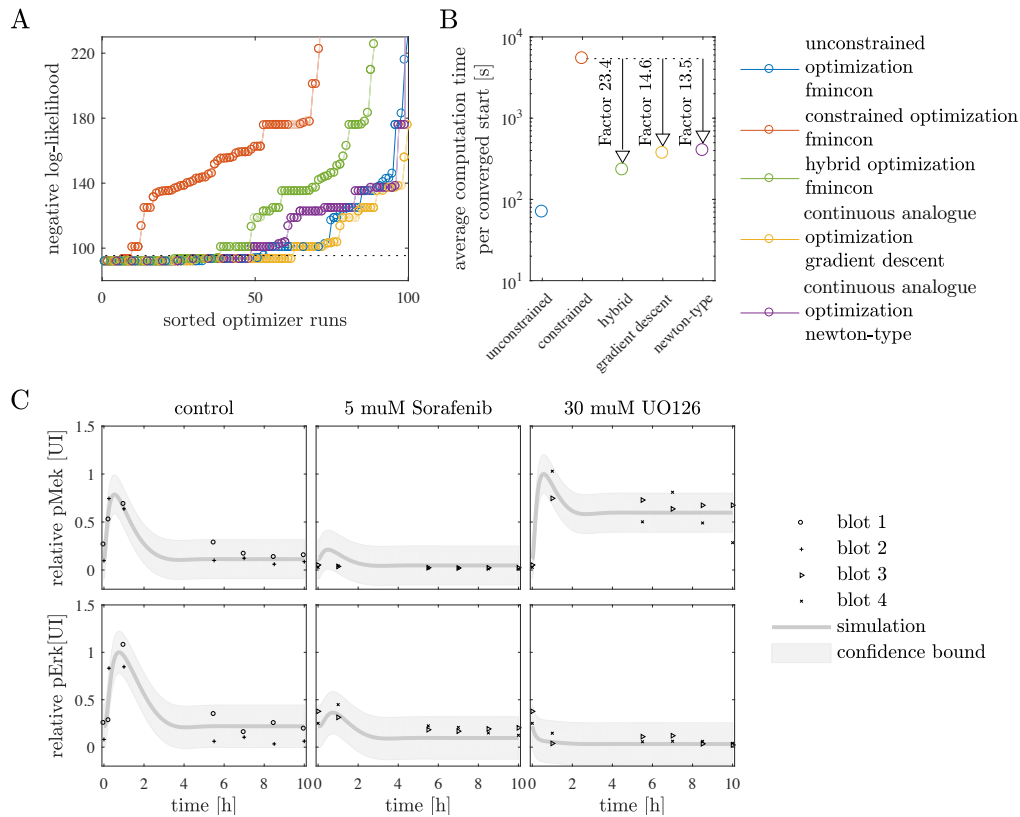


Figure 3.11.: Parameter estimation results for Raf/MEK/ERK signaling in HeLa cells. (A) Convergence and (B) computational efficiency of local optimization methods for the model with the negative feedback loop (H1). (C) Best fit of the model with the negative feedback loop (H1) to data for three different treatment conditions. pMEK and pERK signals are rescaled with the respective maximum activity and the light gray areas indicates $2\text{-}\sigma$ interval of the measurement noise.

The starting points for local optimizations were obtained using Latin hypercube sampling (see Table 3.2). The maximal number of iterations and function evaluations performed by `fmincon` were increased to 2000 and $2000n_\theta$ for the unconstrained and constrained optimization. For the hybrid optimization, the maximal number of iterations was increased to 2000. The results for 100 starts of the local optimizations for the model of H1 are depicted in Figure 3.11 A and B.

3.3.4.4. Hybrid and continuous analogue optimization outperforms constrained optimization

Unconstrained optimization using the analytical expression for the steady state – the gold standard – converged in $\sim 50\%$ of the starts (Figure 3.11 A). Hybrid and continuous

analogue optimization methods achieved a percentage of converged starts comparable to the gold standard (40-60%), but without requiring an analytical expression for the steady state. Constrained optimization – the state-of-the-art – converged in less than 10% of the starts, resulting in a relatively large computation time per converged start (Figure 3.11 B). Even though hybrid and continuous analogue optimization were slower than the gold standard, they were more than 10 times faster than constrained optimization. Hence, the proposed optimization methods also outperformed constrained optimization for this problem.

A detailed comparison of the proposed methods revealed that continuous analogue optimization using gradient descent achieved the highest percentage of converged starts. However, hybrid optimization required fewer simulations of the perturbation experiments – the time-consuming step – rendering this method computationally more efficient. Continuous analogue optimization using Newton-type descent was the least efficient of the proposed methods. This might be related to the challenges in tuning the regularization parameters.

3.3.4.5. Model selection reveals importance of negative feedback

The model with negative feedback (H1) fitted the experimental data (Figure 3.11 C). It captured the transient phosphorylation of MEK and ERK after release from S-phase arrest and the reduced ERK phosphorylation in the presence of Sorafenib and UO126. Furthermore, the increased MEK phosphorylation after UO126 treatment is explained via a decrease in the strength of the negative feedback which is caused by the reduced abundance of pERK. The model without the negative feedback loop (H2) was not able to capture the difference between the control condition and the simulation with UO126. The value of the Bayesian Information Criterion (BIC) (Schwarz, 1978) is 278.4 for the model with negative feedback (H1) and 317.4 for the model without negative feedback (H2). The difference of 39.0 indicates a strong preference for H1 (Burnham and Anderson, 2002). The same conclusion was reached using the Akaike Information Criterion (AIC) (Akaike, 1978). We conclude that Raf phosphorylation is inhibited by pERK during G1/S phase transition.

To summarize, in this section we illustrated the proposed hybrid and continuous analogue optimization methods. The applicability of the methods was demonstrated by studying relevant biological problems. The comparison with state-of-the-art methods revealed convergence and computational efficiency. The study of Raf/MEK/ERK signaling using the methods underlined the feedback regulation of ERK phosphorylation during cell cycle progression.

3.4. Discussion

Optimization problems with steady-state constraints arise in many biology applications for a wide range of ODE models (Hass et al., 2019). For some models an analytical expression for the steady state can be derived and used to eliminate the steady-state constraints (Rosenblatt et al., 2016). While this is favorable, it is not always possible. In cases in which no analytical expressions are available, the vector of optimization variables contains the unknown parameters as well as the corresponding steady states. The optimizers have to evolve on the non-linear manifold, the set of steady states. To address Problem 3.1.1 posed in Section 3.1.4, we proposed a hybrid optimization method and a continuous analogue to solve optimization problems with steady-state constraints more efficiently by exploiting the local geometry of the steady-state manifold. For the continuous analogue, we established local asymptotic stability of optimal points using singular perturbation theory. This result is however restricted to locally strictly convex objective functions implying local practical identifiability. The properties in the presence of practical and structural non-identifiability remain to be analyzed. Preliminary results and the applications suggest that in the presence of non-identifiabilities the continuous analogue optimization method yields a point on the non-identifiable subspace. Furthermore, the available proof shows the retraction factor λ has to be chosen large enough to ensure convergence. However, as too large λ will result in a stiff system, an intelligent choice of λ is necessary.

The proposed hybrid and continuous analogue optimization methods were evaluated using five models for biological processes. Following simple illustration examples using simulated data, an application with multiple steady-state constraints and an application with time-resolved data for multiple perturbation conditions were considered. For this rich set of scenarios we found that the hybrid and the continuous analogue optimization methods possessed improved convergence properties in comparison to standard constrained optimization methods implemented in the MATLAB routine `fmincon`. We expect that the proposed methods also outperform alternative optimization routines, e.g., IPOPT (Wächter and Biegler, 2006). However, this remains to be analyzed. The proposed optimization methods yielded convergence properties comparable to those of unconstrained optimization methods exploiting an analytical expression for the steady state. However, if analytical expressions for the steady state can be determined using available methods (Halász et al., 2013; Loriaux et al., 2013; Rosenblatt et al., 2016), unconstrained optimization should be used as the computation time is lower. The proposed methods are also applicable to a broader class of problems for which no analytical expression for the steady

state is available. Furthermore, the method directly allows for multiple steady-state constraints. Unlike methods based on sequential geometric programming (Pozo et al., 2011; Xu, 2013), steady-state and kinetic data can be incorporated.

The implementation of the hybrid optimization method employed in this study is a simulation-based retraction operator. Alternatively, efficient and accurate schemes combining simulation and local optimization could be employed to compute steady states and sensitivities (Shiraishi et al., 2014). This should improve the computational efficiency further.

Preliminary results suggested that the methods also achieve good convergence for dynamical systems with multiple stable steady states and bifurcations (Müller and Kuttler, 2015) (Section 3.3.2). The theoretical analysis of the proposed methods and a detailed performance evaluation for dynamical systems with such properties remain to be addressed.

In summary, the proposed optimization methods are promising alternatives to constrained optimization for optimization problems with steady-state constraints. They are applicable to a wide range of ODE-constrained optimization problems (Hasenauer et al., 2014; Zechner et al., 2012) and can – unlike methods which rely on an analytical expression for the steady state – be extended to PDE constrained optimization problems (Hock et al., 2013).

Chapter 4.

PDE constrained problems with additional steady-state constraints

In the previous chapter, we developed, among others, a continuous analogue for the optimization of ODE constrained problems with an equality constraint derived from a steady-state condition. A natural extension is to consider biological processes that are described by PDEs. In this chapter, we introduce the mathematical formulation of the research question (Section 4.1). Consequently, we extend the continuous analogue framework to PDE constrained optimization problems (Section 4.2) and analyze the stability of optimal points (Section 4.3). We illustrate the novel approach using simulated data of a model of gradient formation in biological tissues and compare the novel approach to standard optimization (Section 4.4). Beyond the generalization of the results of the previous chapter for ODEs to function spaces, we provide rationales and constraints for the choice of tuning parameters, in particular of the retraction factor. The properties of the continuous analogues are studied in for a model of gradient formation in biological tissues.

The content, text, data and figures of this chapter are based on and in part identically with the paper *Continuous analogue to iterative optimization for PDE-constrained inverse problems* by Romana Boiger, Anna Fiedler, Jan Hasenauer and Barbara Kaltenbacher published in *Inverse Problems in Science and Engineering*, 27:6, 710–734, 2019 (Boiger et al., 2019).

4.1. Introduction and mathematical problem formulation

In the field of mathematical biology, PDEs are often used to describe spatio-temporal processes (Saunders et al., 2012), population dynamics (Luzyanina et al., 2014) and stochastic

processes (Risken and Frank, 1996). In some of these applications steady states arise naturally due to fast equilibration. For example substance gradients are formed and sustained in some tissues to provide positional information to cells sensing the substance (see also Chapter 4.4). Further, a process might also start initially in steady state and then show transient dynamic behavior caused by the change of some external factors. Both cases lead to additional steady-state constraints for the parameter estimation. In the following, we introduce steady-state constraints for PDE constrained optimization for elliptic and parabolic PDEs and formulate the research problem.

4.1.1. Formulation of mathematical model

In this section, we describe the general setting for PDE constrained optimization with additional steady-state constraints. To embed the PDE models in a functional analytic setting, we consider a separable Banach space V with dual space V^* . Further, let V be continuously and densely embedded into a Hilbert space H , such that $V \subseteq H \cong H^* \subseteq V^*$ forms a Gelfand triple.

We study elliptic and parabolic PDE models. In the parabolic case, the initial condition $u_0 \in V$ of the parabolic PDE is defined as the solution of an elliptic PDE,

$$\begin{aligned} u_t(t) &= C(\theta, t, u(t)), \quad t \in]0, T[\\ u(0) &= u_0, \end{aligned} \tag{4.1}$$

with

$$0 = C_0(\theta, u_0), \tag{4.2}$$

in which u_0 denotes a stable steady state of the unperturbed system, while $u \in W(0, T) = L^2(0, T; V) \cap H^1(0, T; V^*)$ denotes the transient solution of the perturbed system starting in the steady state of the unperturbed system u_0 . The parameters θ are assumed to be finite dimensional and real-valued, $\theta \in \mathbb{R}^{n_\theta}$ and the dynamics are given by the operators $C: \mathbb{R}^{n_\theta} \times]0, T[\times V \rightarrow V^*$ and $C_0: \mathbb{R}^{n_\theta} \times V \rightarrow V^*$. For this PDE, existence of a weak solution holds, e.g., under Assumption 2.1.1 (Zeidler, 1990, p.770 ff.) for C . Here, we assume that the assumption holds for all parameters θ in a neighborhood around the optimal parameters θ^* , i.e., $\theta \in B_{\varepsilon_\theta}(\theta^*) = \{\theta \in \mathbb{R}^{n_\theta} \mid \|\theta - \theta^*\| < \varepsilon_\theta\} \subseteq \mathbb{R}^{n_\theta}$ for some $\varepsilon_\theta > 0$. For the differential operator C_0 , the first two items in Assumption 2.1.1 are assumed to hold, cf. Assumption 4.3.2 below.

For each t and θ the observables of the models are defined via observation operators $B: \mathbb{R}^{n_\theta} \times]0, T[\times V \rightarrow Z$ and $B_0: \mathbb{R}^{n_\theta} \times V \rightarrow Z_0$ for some observation spaces Z, Z_0 that

4.1. Introduction and mathematical problem formulation

are Hilbert spaces,

$$y(\theta, t) = B(\theta, t, u(t)), \quad y_0(\theta) = B_0(\theta, u_0). \quad (4.3)$$

In mathematical biology, the differential operator C is often semi-linear and describes a reaction-diffusion-advection equation,

$$C(\theta, t, u) = f(k, u) - \nabla_x \cdot (vu - D\nabla_x u),$$

in which $u(t) \in V$ is a concentration vector, $x \in \Omega \subseteq \mathbb{R}^n$ is the spatial location, and $f: \mathbb{R}^{n_k} \times V \rightarrow V^*$ is the reaction term. The parameters $\theta = (k, v, D)$ are the velocity vector $v \in \mathbb{R}^n$, the diffusion matrix $D \in \mathbb{R}^{n \times n}$, and the kinetic parameters $k \in \mathbb{R}^{n_k}$. Assumption 2.1.1 is satisfied, if D is positive definite and f fulfills a certain growth condition.

4.1.2. Parameter estimation problem

The unknown model parameters θ are estimated from noise-corrupted measurements of the observables y by minimizing an objective function, e.g., the negative log-likelihood or the sum-of-squared-residuals. In the following, we distinguish two cases: Problems with elliptic and parabolic constraints and problems with only elliptic constraints.

Elliptic and parabolic PDE constraints In the general case, observations are available for the initial state and the transient phase. The objective function \tilde{J} depends on the parameters and the parameter-dependent solutions of the parabolic and the elliptic PDE, $\tilde{J}: \mathbb{R}^{n_\theta} \times V \times W(0, T) \rightarrow \mathbb{R}$, e.g., $\tilde{J}(\theta, u_0, u) = \frac{1}{2} \|y_0 - B_0(\theta, u_0)\|_{Z_0}^2 + \frac{1}{2} \int_0^T \|y(t) - B(\theta, t, u)\|_Z^2 dt$. The optimization problem is given by

$$\begin{aligned} \min_{\theta, u_0, u} \quad & \tilde{J}(\theta, u_0, u) \\ \text{s.t.} \quad & \begin{cases} u_t = C(\theta, t, u), & u(0) = u_0 \\ 0 = C_0(\theta, u_0) \end{cases} \end{aligned} \quad (4.4)$$

Elliptic PDE constraint In many applications, only experimental data for the steady state of a process are available, e.g., due to fast equilibration of the process. In this case, the problem is simplified and the objective function J depends only on the parameters

and parameter-dependent solutions of the elliptic PDE, $J: \mathbb{R}^{n_\theta} \times V \rightarrow \mathbb{R}$, e.g., $J(\theta, u_0) = \frac{1}{2} \|y_0 - B_0(\theta, u_0)\|_{Z_0}^2$. The optimization problem is given by

$$\begin{aligned} \min_{\theta, u_0} J(\theta, u_0) \\ \text{s.t. } 0 = C_0(\theta, u_0). \end{aligned} \tag{4.5}$$

The reduced formulation of the optimization problem (4.5) is given by

$$\min_{\theta} \{j(\theta) := J(\theta, \varphi_0(\theta))\}, \tag{4.6}$$

in which $\varphi_0(\theta) \in V$ now denotes the (infinite dimensional) parameter-dependent solution of $C_0(\theta, u_0) = 0$ and $j: \mathbb{R}^{n_\theta} \rightarrow \mathbb{R}$ denotes the reduced objective function.

4.1.3. Problem formulation

Most numerical methods for PDE-constrained optimization problems and literature on parameter estimation in special applications focus on iterative methods, which generate a discrete sequence of points along which the objective function decreases (Banks and Kunsch (1989); Bock et al. (2013); Carvalho et al. (2015); Hinze et al. (2009); Ito and Kunisch (2008); Nielsen et al. (2013); Xun et al. (2013) and references therein). However, for many constrained and unconstrained optimization problems, continuous analogues exhibit larger regions of attraction and more robust convergence than discrete iterative methods (Tanabe, 1985). Based on the findings of improved convergence and run time in the previous chapter, we set out to develop a continuous analogue approach for PDE constrained optimization problems (see Airapetyan et al. (2000); Kaltenbacher et al. (2002); Tanabe (1979, 1980, 1985); Watson (2001) and references therein).

Problem 4.1.1. *Given optimization problems (4.4) and (4.5), develop a continuous analogue for parameter estimation that exploits the local structure of the problem and investigate the properties of the developed approach.*

As already outlined in the previous chapter, the continuous analogue depends on a tuning parameter, the retraction factor λ . It is not clear how to choose the retraction factor a priori. Too small choices might not guarantee convergence and too large choices increase the stiffness of the continuous analogue. Hence, we also consider the problem:

Problem 4.1.2. *Given the continuous analogue, determine the influence of the retraction factor and derive a lower bound for it.*

4.2. Continuous analogue of descent methods for PDE constrained problem

In this section, we develop a continuous analogue of an iterative descent method for optimization problems of type (4.4) and (4.5). For simplicity, we first consider elliptic PDE constraints (4.5) and afterwards generalize the results to mixed parabolic and elliptic PDE constraints (4.4).

4.2.1. Elliptic PDE constraints

To address Problem 4.1.1, we derive a coupled ODE-PDE system for the solution of optimization problem (4.5). The trajectory of this continuous analogue evolves in parameter and state space on the manifold defined by $C_0(\theta, u_0) = 0$ towards a local minimum by exploiting the first order geometry of the manifold, i.e., its tangent space.

Mathematically, the first order geometry of $C_0(\theta, \varphi_0(\theta)) = 0$ is defined by the sensitivity equations

$$\frac{\partial C_0}{\partial u_0}(\theta, \varphi_0(\theta)) \frac{\partial \varphi_0}{\partial \theta_i}(\theta) + \frac{\partial C_0}{\partial \theta_i}(\theta, \varphi_0(\theta)) = 0, \quad i \in \{1, \dots, n_\theta\}. \quad (4.7)$$

The sensitivity equations can be reformulated to

$$\frac{\partial \varphi_0}{\partial \theta_i}(\theta) = - \left(\frac{\partial C_0}{\partial u_0}(\theta, \varphi_0(\theta)) \right)^{-1} \frac{\partial C_0}{\partial \theta_i}(\theta, \varphi_0(\theta)), \quad i \in \{1, \dots, n_\theta\},$$

provided the inverse of $\frac{\partial C_0}{\partial u_0}(\theta, \varphi_0(\theta))$ exists. We extend $\frac{\partial \varphi_0}{\partial \theta_i}$ to points (θ, u_0) not necessarily lying on the solution manifold of $C_0 = 0$ by defining the operator $S_0: \mathbb{R}^{n_\theta} \times V \rightarrow V^{n_\theta}$, that provides the solution of the sensitivity equations

$$\frac{\partial C_0}{\partial u_0}(\theta, u_0) S_{0i}(\theta, u_0) + \frac{\partial C_0}{\partial \theta_i}(\theta, u_0) = 0, \quad i \in \{1, \dots, n_\theta\}, \quad (4.8)$$

for given θ and u_0 . With (4.7) it holds that

$$S_0(\theta, \varphi_0(\theta)) = \nabla_\theta \varphi_0(\theta), \quad (4.9)$$

provided the solutions to (4.7) and (4.8) are unique.

In order to couple changes in θ with appropriate changes in the state variable u_0 , we use the fact that the function $\nabla_\theta \varphi_0(\theta)$ provides the first order term of the Taylor series

expansion of the steady state with respect to the parameter vector θ ,

$$\varphi_0(\theta + r\Delta\theta) = \varphi_0(\theta) + \nabla_{\theta}\varphi_0(\theta)r\Delta\theta + o(r) \quad \text{as } r \rightarrow 0, r \in \mathbb{R}. \quad (4.10)$$

Defining $\hat{\varphi}_0(r) := \varphi_0(\theta + r\Delta\theta)$ for some $\Delta\theta \in \mathbb{R}^{n_{\theta}}$ and differentiating (4.10) with respect to r yields

$$\frac{d\hat{\varphi}_0}{dr}(r) = \nabla_{\theta}\varphi_0(\theta)\Delta\theta + o(1) = S_0(\theta + r\Delta\theta, \hat{\varphi}_0(r))\Delta\theta + o(1) \quad \text{as } r \rightarrow 0. \quad (4.11)$$

This relation motivates the formulation of the coupled ODE-PDE model in θ and u_0

$$\begin{aligned} \frac{d\theta}{dr}(r) &= g(\theta, u_0), & \theta(0) &= \theta_0 \\ \frac{du_0}{dr}(r) &= S_0(\theta, u_0)\frac{d\theta}{dr}(r) = S_0(\theta, u_0)g(\theta, u_0), & u_0(0) &= u_{0,0} \end{aligned} \quad (4.12)$$

using the artificial time parameter r . For a change in the parameters $\frac{d\theta}{dr}$, the update in u_0 is chosen according to (4.11). Solutions of this dynamical system evolve on the manifold $C_0(\theta(r), u_0(r)) = 0$ for arbitrary parameter update directions $g: \mathbb{R}^{n_{\theta}} \times V \rightarrow \mathbb{R}^{n_{\theta}}$, provided that the initial state is on the manifold $C_0(\theta_0, u_{0,0}) = 0$. The state variables of this coupled ODE-PDE system are θ and u_0 , and the path variable is r . To solve optimization problem (4.6), g is chosen as an arbitrary descent direction satisfying

$$\nabla j(\theta)^T g(\theta, u_0) < 0, \quad (4.13)$$

more precisely satisfying Assumption 4.3.5 below. For example, g can be chosen as a steepest descent direction

$$g = \underset{\|v\|_* \leq 1}{\operatorname{argmin}} \nabla j(\theta)^T v \quad (4.14)$$

for some norm $\|\cdot\|_*$. For the Euclidian norm we obtain the gradient descent direction,

$$g_i(\theta, u_0) := -\frac{\partial J}{\partial \theta_i}(\theta, u_0) - \left\langle \frac{\partial J}{\partial u_0}(\theta, u_0), S_{0i}(\theta, u_0) \right\rangle_{V^*, V} =: d_i(\theta, u_0), \quad i \in \{1, \dots, n_{\theta}\}, \quad (4.15)$$

in which we substituted $\varphi_0(\theta)$ by u_0 extending the definition also to states u_0 that are not on the steady-state manifold. Likewise, defining $\|v\|_*^2 := v^T H(\theta, u_0)v$ with some positive definite matrix $H(\theta, u_0)$, so that

$$g(\theta, u_0) = H(\theta, u_0)^{-1}d(\theta, u_0) \quad (4.16)$$

leads to a descent direction. Using, e.g., the Hessian of j leads to Newton, Gauss-Newton (by excluding second order sensitivities) or quasi-Newton methods (where approximations to the Hessian are computed via low rank updates). As the Hessian is not guaranteed to be positive definite, regularization with a scaled identity matrix, $H_\mu(\theta, u_0) = H(\theta, u_0) + \mu I$ with $\mu > 0$, might be useful. However, how to choose μ and potential continuous update rules are out of the scope of this work.

The coupled ODE-PDE systems (4.12) can be solved using numerical time stepping methods. These numerical methods might however accumulate errors resulting in the divergence of the state $(\theta(r), u_0(r))$ from the steady-state manifold. Additionally, the initial state, $u_{0,0}$ might not be on the steady-state manifold. To account for this, we include the retraction term $\lambda C_0(\theta, u_0)$ in the evolution equation of u_0 , with retraction factor $\lambda > 0$. This yields the following continuous analogue of a descent method for optimization problems with elliptic PDE constraints,

$$\begin{aligned} \frac{d\theta}{dr}(r) &= g(\theta, u_0), & \theta(0) &= \theta_0 \\ \frac{du_0}{dr}(r) &= S_0(\theta, u_0)g(\theta, u_0) + \lambda C_0(\theta, u_0), & u_0(0) &= u_{0,0}. \end{aligned} \quad (4.17)$$

As, for fixed θ , the equation $C_0(\theta, u_0) = 0$ defines a stable steady state of the PDE (4.2), the retraction term stabilizes the manifold. For $\lambda \gg 1$, the system should first converge to the steady state $\varphi_0(\theta_0)$ for the initial parameter θ_0 and then move along the manifold to a local optimum θ^* as illustrated in Figure 4.1.

4.2.2. Elliptic and parabolic PDE constraints

The continuous analogue for descent methods with elliptic PDE constraints can be generalized to problems with parabolic and elliptic PDE constraints. One possibility for doing so is to consider the partially reduced problem

$$\begin{aligned} \min_{\theta, u_0} \tilde{\mathcal{J}}(\theta, u_0) &:= \tilde{J}(\theta, u_0, \varphi(\theta, u_0)) \\ \text{s.t. } 0 &= C_0(\theta, u_0). \end{aligned} \quad (4.18)$$

in which $u = \varphi(\theta, u_0)$ denotes the solution to $u_t = C(\theta, t, u)$ with $u(0) = u_0$. Given this formulation, we can use continuous analogue (4.17) with

$$\tilde{g}_i(\theta, u_0) = -\frac{\partial \tilde{\mathcal{J}}}{\partial \theta_i}(\theta, u_0) - \frac{\partial \tilde{\mathcal{J}}}{\partial u_0}(\theta, u_0) S_{0i}(\theta, u_0), \quad i \in \{1, \dots, n_\theta\}. \quad (4.19)$$

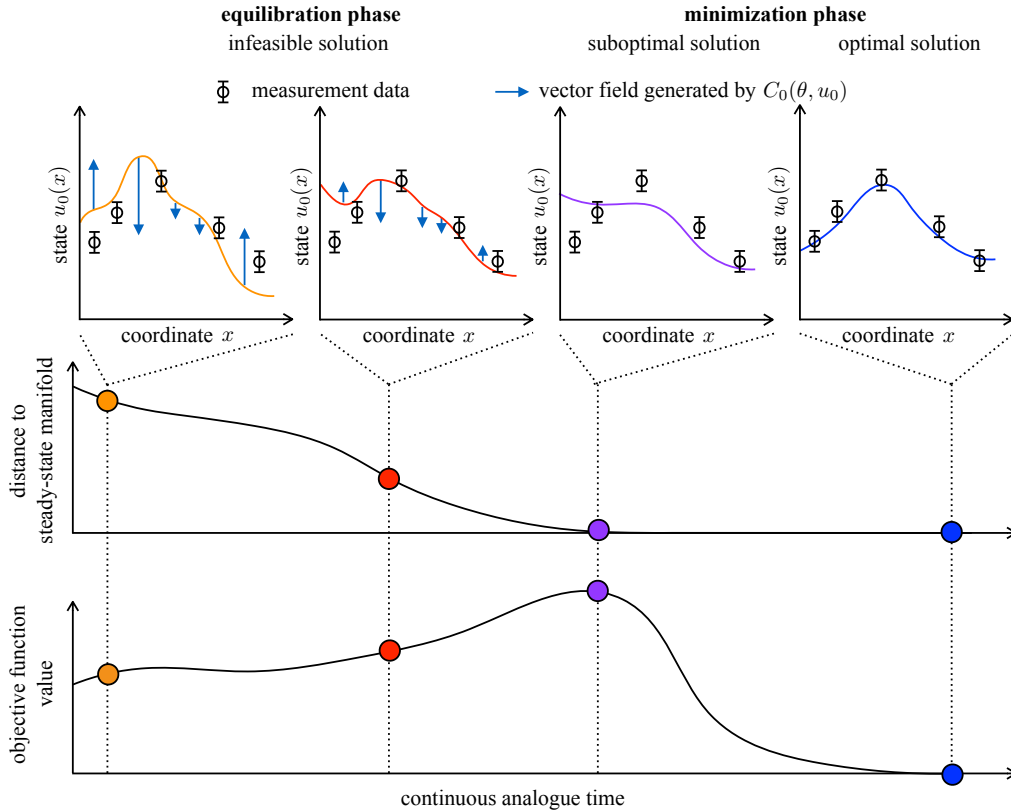


Figure 4.1.: The state of the system is illustrated along the trajectory of (4.17). In the first phase, the equilibration phase, the system converges to the manifold. The solution is not feasible during this phase as the equality constraint, $C_0(\theta, u_0) = 0$, is violated. In the course of the equilibration the objective function value might increase. In the second phase, the minimization phase, the objective function is minimized along the steady-state manifold.

To avoid the need for the solution operator $\varphi: \mathbb{R}^{n_\theta} \times V \rightarrow W(0, T)$, alternatively, a continuous analogue of the full problem can also be formulated. This is beyond the scope of this work and will be subject of future research.

4.3. Local stability and convergence to a local optimum

The behavior of the coupled ODE-PDE system (4.17) introduced in the previous section depends on the properties of the objective function and the PDE model, as well as the retraction factor λ . To prove that a solution of (4.17) with an appropriate retraction factor λ is well defined and converges to the local minimizer $(\theta^*, u_0^*) = (\theta^*, \varphi_0(\theta^*))$ of (4.5), we impose the following assumptions:

4.3. Local stability and convergence to a local optimum

Assumption 4.3.1. *The descent direction vanishes at the minimizer θ^* of the optimization problem $\min_{\theta} j(\theta)$,*

$$g(\theta^*, \varphi_0(\theta^*)) = 0.$$

Assumption 4.3.2. *There exists $\varepsilon_{\theta} > 0$ such that for all $\theta \in B_{\varepsilon_{\theta}}(\theta^*) = \{\theta \in \mathbb{R}^{n_{\theta}} \mid \|\theta - \theta^*\| < \varepsilon_{\theta}\} \subseteq \mathbb{R}^{n_{\theta}}$, the following holds:*

- *The operator $-C_0(\theta, \cdot)$ is monotone and hemicontinuous.*
- *The operator $-C_0(\theta, \cdot)$ is coercive.*

Assumption 4.3.3. *The function $C_0(\theta, u_0)$ is locally uniformly monotonically decreasing, i.e., there exist $\gamma_c > 0$ and $\varepsilon_{u_0} > 0$ such that*

$$\langle C_0(\theta, u_0^1) - C_0(\theta, u_0^2), u_0^1 - u_0^2 \rangle_{V^*, V} \leq -\gamma_c \|u_0^1 - u_0^2\|_V^2$$

for all $\theta \in B_{\varepsilon_{\theta}}(\theta^*)$ and $u_0^1, u_0^2 \in B_{\varepsilon_{u_0}}(\varphi_0(\theta^*)) := \{u_0 \in V \mid \|u_0 - \varphi_0(\theta^*)\|_V < \varepsilon_{u_0}\}$.

Assumption 4.3.4. *The sensitivity $S_0(\theta, u_0)$ is locally Lipschitz continuous with respect to u_0 , i.e., there exists $L_{S_0} \geq 0$ such that*

$$\|S_0(\theta, u_0^1) - S_0(\theta, u_0^2)\|_{V^{n_{\theta}}} \leq L_{S_0} \|u_0^1 - u_0^2\|_V$$

for all $\theta \in B_{\varepsilon_{\theta}}(\theta^*)$ and $u_0^1, u_0^2 \in B_{\varepsilon_{u_0}}(\varphi_0(\theta^*))$.

Assumption 4.3.5. *The mapping $\theta \mapsto g(\theta, \varphi_0(\theta))$ is uniformly monotonically decreasing on $B_{\varepsilon_{\theta}}(\theta^*)$, i.e., there exists $\gamma_g > 0$ such that*

$$(g(\theta, \varphi_0(\theta)) - g(\theta^*, \varphi_0(\theta^*)))^T (\theta - \theta^*) \leq -\gamma_g \|\theta - \theta^*\|^2$$

for all $\theta \in B_{\varepsilon_{\theta}}(\theta^*)$.

Assumption 4.3.6. *The descent direction g is locally Lipschitz continuous with respect to u_0 , i.e., there exists $L_g \geq 0$ such that*

$$\|g(\theta, u_0^1) - g(\theta, u_0^2)\| \leq L_g \|u_0^1 - u_0^2\|_V$$

for all $\theta \in B_{\varepsilon_{\theta}}(\theta^*)$ and $u_0^1, u_0^2 \in B_{\varepsilon_{u_0}}(\varphi_0(\theta^*)) \subseteq V$.

Moreover, g is uniformly bounded on $B_{\varepsilon_{\theta}}(\theta^*) \times B_{\varepsilon_{u_0}}(\varphi_0(\theta^*))$, i.e there exists $K_g \geq 0$ such that

$$\|g(\theta, u_0)\| \leq K_g$$

for all $(\theta, u_0) \in B_{\varepsilon_{\theta}}(\theta^*) \times B_{\varepsilon_{u_0}}(\varphi_0(\theta^*))$.

4.3.1. Elliptic PDE constraints

Using Assumptions 4.3.1-4.3.6 and the existence of a weak solution (Assumption 2.1.1), we can prove the following theorem on stability and convergence for the continuous analogue of the descent method for elliptic PDE constraints:

Theorem 4.3.1. *Let Assumptions 4.3.1-4.3.6 be satisfied. Then there exists a $\lambda^* \geq 0$ such that for all $\lambda > \lambda^*$ solutions to (4.17) are well-defined for all $r > 0$ and the local minimizer (θ^*, u_0^*) of the optimization problem (4.5) is a locally exponentially stable steady state of the system (4.17).*

Proof. Define $\tilde{\theta} := \theta - \theta^*$ and $\tilde{u}_0 := u_0 - \varphi_0(\theta)$, with $\theta \in B_{\varepsilon_\theta}(\theta^*)$ and $u_0 \in B_{\varepsilon_{u_0}}(\varphi_0(\theta^*))$, where $\varphi_0(\theta^*) = u_0^*$ exists, because of Assumption 4.3.2. We further define a Lyapunov function $\mathcal{V}(r) = \frac{1}{2}\|\tilde{\theta}(r)\|^2 + \frac{1}{2}\|\tilde{u}_0(r)\|_H^2$. To prove Theorem 4.3.1, we will show that the Lyapunov function decreases exponentially. The derivative along the trajectories is given by

$$\frac{d}{dr}\mathcal{V}(r) = \frac{d}{dr}\left(\frac{1}{2}\|\tilde{\theta}(r)\|^2\right) + \frac{d}{dr}\left(\frac{1}{2}\|\tilde{u}_0(r)\|_H^2\right).$$

First, we bound the first summand from above, using (4.17), and Assumptions 4.3.1, 4.3.5 and 4.3.6,

$$\begin{aligned} \frac{d}{dr}\frac{1}{2}\|\tilde{\theta}\|^2 &= \left(\frac{d}{dr}\tilde{\theta}\right)^T \tilde{\theta} \\ &= (g(\theta, u_0) - g(\theta, \varphi_0(\theta)))^T \tilde{\theta} + (g(\theta, \varphi_0(\theta)) - g(\theta^*, \varphi_0(\theta^*)))^T \tilde{\theta} \\ &\leq \|g(\theta, u_0) - g(\theta, \varphi_0(\theta))\| \|\tilde{\theta}\| - \gamma_g \|\theta - \theta^*\|^2 \\ &\leq L_g \|\tilde{u}_0\| \|\tilde{\theta}\| - \gamma_g \|\tilde{\theta}\|^2. \end{aligned}$$

Second, we bound the second summand from above, using (4.17) and $C_0(\theta, \varphi_0(\theta)) = 0$, as well as the fact that by Assumption 4.3.3 we have (4.9),

$$\begin{aligned} \frac{d}{dr}\frac{1}{2}\|\tilde{u}_0\|_H^2 &= \left(\frac{d\tilde{u}_0}{dr}, \tilde{u}_0\right)_{V^*, V} = \left(\frac{du_0}{dr} - \nabla_{\theta}\varphi_0(\theta)\frac{d\theta}{dr}, \tilde{u}_0\right)_{V^*, V} \\ &= ((S_0(\theta, u_0) - S_0(\theta, \varphi_0(\theta)))g(\theta, u_0), \tilde{u}_0)_{V^*, V} \\ &\quad + \lambda(C_0(\theta, u_0) - C_0(\theta, \varphi_0(\theta)), \tilde{u}_0)_{V^*, V} \\ &\leq \|(S_0(\theta, u_0) - S_0(\theta, \varphi_0(\theta)))\|_{V^*, V} \|g(\theta, u_0)\| \|\tilde{u}_0\|_V \\ &\quad + \lambda(C_0(\theta, u_0) - C_0(\theta, \varphi_0(\theta)), \tilde{u}_0)_{V^*, V}. \end{aligned}$$

4.3. Local stability and convergence to a local optimum

With Assumptions 4.3.3, 4.3.4 and 4.3.6 we get

$$\frac{d}{dr} \frac{1}{2} \|\tilde{u}_0\|_H^2 \leq (L_{S_0} K_g - \lambda \gamma_c) \|\tilde{u}_0\|_V^2.$$

Hence, we can estimate the derivative of the Lyapunov function,

$$\frac{d}{dr} \mathcal{V}(r) \leq -(-L_{S_0} K_g + \lambda \gamma_c) \|\tilde{u}_0\|_V^2 + L_g \|\tilde{u}_0\|_V \|\tilde{\theta}\| - \gamma_g \|\tilde{\theta}\|^2.$$

To show that \mathcal{V} decays exponentially we have to show that

$$\frac{d}{dr} \mathcal{V}(r) \leq -a \mathcal{V}(r)$$

for some $a > 0$. Based on our estimates, proving Theorem 4.3.1 reduces to finding $a > 0$ with

$$0 \leq (-L_{S_0} K_g + \lambda \gamma_c - \frac{a}{2}) \|\tilde{u}_0\|_V^2 - L_g \|\tilde{u}_0\|_V \|\tilde{\theta}\| + (\gamma_g - \frac{a}{2}) \|\tilde{\theta}\|^2 \quad (4.20)$$

We want this inequality to be valid without restrictions on $\|\tilde{\theta}\|$ or $\|\tilde{u}_0\|_V$. Due to the last term, we can therefore only consider values of a that are smaller than $2\gamma_g$. Hence, (4.20) is equivalent to

$$0 \leq \left(\sqrt{\gamma_g - \frac{a}{2}} \|\tilde{\theta}\| - \frac{L_g}{2\sqrt{\gamma_g - \frac{a}{2}}} \|\tilde{u}_0\|_V \right)^2 + \left(-L_{S_0} K_g + \lambda \gamma_c - \frac{a}{2} - \frac{L_g^2}{4(\gamma_g - \frac{a}{2})} \right) \|\tilde{u}_0\|_V^2.$$

Since the first term in the inequality is greater or equal to 0, we have to find $a > 0$ such that

$$\lambda \gamma_c - \frac{a}{2} - L_{S_0} K_g - \frac{L_g^2}{4(\gamma_g - \frac{a}{2})} \geq 0.$$

Multiplying with $4\gamma_g - 2a$ we obtain a quadratic inequality for a

$$a^2 + 2(-\gamma_g - \lambda \gamma_c + L_{S_0} K_g) a + (4\lambda \gamma_c \gamma_g - 4L_{S_0} K_g \gamma_g - L_g^2) \geq 0.$$

The roots of the quadratic polynomial are given by

$$a_{1,2} = \gamma_g + \lambda \gamma_c - L_{S_0} K_g \pm \sqrt{d},$$

with discriminant

$$d = (\gamma_g - \lambda \gamma_c + L_{S_0} K_g)^2 + L_g^2.$$




Figure 4.2.: The function $f(a) = a^2 + (-2\gamma_g - 2\lambda\gamma_c + 2L_{S_0}K_g)a + (4\lambda\gamma_c\gamma_g - 4L_{S_0}K_g\gamma_g - L_g^2)$ is illustrated with the two roots a_1 and a_2 and the three different positions of $2\gamma_g$, as well as possible positions of a .

The discriminant is always positive, therefore, $a_1 = \gamma_g + \lambda\gamma_c - L_{S_0}K_g - \sqrt{d} < a_2 = \gamma_g + \lambda\gamma_c - L_{S_0}K_g + \sqrt{d}$ are real roots. In the following, we will assume that $a_1 > 0$, which can be achieved by choosing λ such that $\lambda > \lambda^* = \frac{L_{S_0}K_g}{\gamma_c} + \frac{L_g^2}{4\gamma_g\gamma_c} \geq 0$. This choice is justified as follows. As the square root, \sqrt{d} , is always positive, $\gamma_g + \lambda\gamma_c - L_{S_0}K_g > 0$, i.e., $\lambda > \frac{L_{S_0}K_g - \gamma_g}{\gamma_c}$ needs to hold to ensure $a_1 > 0$. Squaring both sides of the inequality

$$\gamma_g + \lambda\gamma_c - L_{S_0}K_g > \sqrt{d} \quad (4.21)$$

yields

$$\begin{aligned} (\gamma_g + \lambda\gamma_c - L_{S_0}K_g)^2 &> (\gamma_g - \lambda\gamma_c + L_{S_0}K_g)^2 + L_g^2 \\ \Leftrightarrow \lambda &> \frac{L_g^2}{4\gamma_g\gamma_c} + \frac{L_{S_0}K_g}{\gamma_c}. \end{aligned}$$

Taking $\lambda > \lambda^* := \max \left\{ \frac{L_{S_0}K_g - \gamma_g}{\gamma_c}, \frac{L_g^2}{4\gamma_g\gamma_c} + \frac{L_{S_0}K_g}{\gamma_c} \right\} = \frac{L_g^2}{4\gamma_g\gamma_c} + \frac{L_{S_0}K_g}{\gamma_c}$ ensures $a_1 > 0$.

Therefore, a either fulfills $0 < a < a_1$ with $a < 2\gamma_g$ or $a_2 < a < 2\gamma_g$, provided $\lambda > \lambda^*$. Hence, we distinguish the following three cases

- (1) $2\gamma_g < a_1 < a_2$,
- (2) $a_1 < a_2 \leq 2\gamma_g$,
- (3) $a_1 \leq 2\gamma_g \leq a_2$;

for the relation of $2\gamma_g$, a_1 and a_2 as illustrated in Figure 4.2.

Case (1): $2\gamma_g < a_1$ is equivalent to

$$\lambda\gamma_c - L_{S_0}K_g - \gamma_g > \sqrt{d}.$$

4.3. Local stability and convergence to a local optimum

If the term $\lambda\gamma_c - L_{S_0}K_g - \gamma_g$ is negative, the inequality cannot be valid. The term $\lambda\gamma_c - L_{S_0}K_g - \gamma_g$ is non-negative if $\lambda \geq \frac{\gamma_g + L_{S_0}K_g}{\gamma_c}$. In this case we can square the inequality and get a contradiction ($0 > L_g^2$).

Case (2): $a_2 \leq 2\gamma_g$ is equivalent to

$$\sqrt{d} \leq \gamma_g + L_{S_0}K_g - \lambda\gamma_c.$$

This leads to a contradiction with the same arguments as in case (1).

Case (3): $a_1 \leq 2\gamma_g$ is equivalent to

$$-\gamma_g + \lambda\gamma_c - L_{S_0}K_g \leq \sqrt{d}.$$

The left hand side $-\gamma_g + \lambda\gamma_c - L_{S_0}K_g$ is non-negative for all $\lambda \geq \frac{\gamma_g + L_{S_0}K_g}{\gamma_c}$. With squaring we get $0 \leq L_g^2$. On the other hand if the term $-\gamma_g + \lambda\gamma_c - L_{S_0}K_g$ is negative, that is $\lambda < \frac{\gamma_g + L_{S_0}K_g}{\gamma_c}$, we have $\sqrt{d} > 0$. This is true for all λ , because $d > 0$. In total, we find $a_1 \leq 2\gamma_g$ for all $\lambda > 0$. Analogously we get for a_2 that $a_2 \geq 2\gamma_g$ is fulfilled for all $\lambda > 0$. Hence, we know that $a_1 \leq 2\gamma_g \leq a_2$ holds for all $\lambda > 0$ and only case (3) is valid.

Altogether, we find that a lies in the interval $[0, a_1]$ provided $\lambda > \lambda^*$. In this case it also holds that

$$\frac{d}{dr} \mathcal{V}(r) \leq -\frac{a}{2} (\|\tilde{u}_0(r)\|_V^2 + \|e(r)\|^2) \leq -\frac{a}{2K_{V \rightarrow H}^2} \|\tilde{u}_0(r)\|_H^2 - \frac{a}{2} \|e(r)\|^2 \leq -\tilde{a} \mathcal{V}(r),$$

with $\tilde{a} = -\frac{a}{2} \min \left\{ \frac{1}{K_{V \rightarrow H}^2}, 1 \right\}$, where $K_{V \rightarrow H}$ is the embedding constant. \square

Remark 4.3.1. *To tune the choice of the retraction factor λ , we now consider the fact that the value of a determines the speed at which $\mathcal{V}(r)$ decreases, thus a convenient choice of the retraction factor $\lambda > \lambda^*$ maximizes a to yield the fastest exponential decay. In our case this means maximizing $a(\lambda) = a_1 = \gamma_g + \lambda\gamma_c - L_{S_0}K_g - \sqrt{d(\lambda)}$ with respect to λ . An elementary computation yields*

$$\frac{da}{d\lambda}(\lambda) = \gamma_c + \frac{\gamma_g\gamma_c - \lambda\gamma_c^2 + L_{S_0}K_g\gamma_c}{\sqrt{d(\lambda)}} \geq 0$$

with equality iff $L_g = 0$, thus $\lambda \mapsto a(\lambda)$ is monotonically increasing (strictly, if $L_g > 0$) and therefore

$$\sup_{\lambda \in (\lambda^*, \infty)} a(\lambda) = \lim_{\lambda \rightarrow \infty} a(\lambda) = \lim_{\lambda \rightarrow \infty} \gamma_g + \lambda\gamma_c - L_{S_0}K_g - \sqrt{(\gamma_g - \lambda\gamma_c + L_{S_0}K_g)^2 + L_g^2} = 2\gamma_g.$$

In case $L_g = 0$ we have $a(\lambda) = \gamma_g + \lambda\gamma_c - L_{S_0}K_g - |\gamma_g - \lambda\gamma_c + L_{S_0}K_g|$. Distinguishing the two cases for the absolute value yields the maximal value $a(\lambda) = 2\gamma_g$, attained at all $\lambda \geq \frac{\gamma_g - L_{S_0}K_g}{\gamma_c}$.

This shows that (unless $L_g = 0$) the exponential decay is maximized by choosing $\lambda > \lambda^*$ as large as possible. Nevertheless, in practice, λ should not be chosen too large in order to avoid stiffness of system (4.17).

Remark 4.3.2. *The proof provides a lower bound for the retraction factor λ , namely $\lambda > \lambda^* = \frac{L_g^2}{4\gamma_g\gamma_c} + \frac{L_{S_0}K_g}{\gamma_c}$ (cf. Problem 4.1.2). In specific applications it might not always be possible to explicitly compute all involved constants. If this is the case, an alternative Lyapunov function can be used to derive a lower bound for λ . A possible candidate for this Lyapunov function is*

$$\mathcal{V}(r) = j(\theta(r)) - j(\theta^*) + \frac{1}{2}\|u_0 - \varphi_0(\theta(r))\|_H^2. \quad (4.22)$$

With this choice and analogous computations as above, different lower bounds involving different constants can be derived. The lower bound for the retraction factor can be estimated as $\lambda > \hat{\lambda}^* = \frac{\hat{L}}{\gamma_c} + \frac{\hat{L}_g^2}{4\gamma_c}$ with

$$\hat{L} = \begin{cases} \frac{((S_0(\theta, u_0) - S_0(\theta, \varphi_0(\theta)))g(\theta, u_0), u_0 - \varphi_0(\theta))_{V^*, V}}{\|u_0 - \varphi_0(\theta)\|_V^2}, & \text{if } u_0 \neq \varphi_0(\theta) \\ 0, & \text{else} \end{cases} \quad (4.23)$$

and

$$\hat{L}_g = \begin{cases} \frac{g(\theta, \varphi_0(\theta))^T(g(\theta, \varphi_0(\theta)) - g(\theta, u_0))}{\|u_0 - \varphi_0(\theta)\|_V \|g(\theta, \varphi_0(\theta))\|}, & \text{if } u_0 \neq \varphi_0(\theta) \text{ and } g(\theta, \varphi_0(\theta)) \neq 0 \\ 0, & \text{else.} \end{cases} \quad (4.24)$$

This bound depends on θ , the current parameter estimates during computation, and therefore requires a posteriori adaptation of the retraction factor. A practical implementation of such a retraction factor choice involves evaluation of functionals of the residuals $g(\theta, \varphi_0(\theta)) - g(\theta, u_0)$, $u_0 - \varphi_0(\theta)$ as well as sensitivities $(S_0(\theta, u_0) - S_0(\theta, \varphi_0(\theta)))g(\theta, u_0)$ (which can be done approximately, on a coarser computational mesh, and using adjoint techniques).

Remark 4.3.3. *If the local minimum at θ^* is not strict and we find a submanifold of optimal parameters, i.e., the parameters are structurally unidentifiable, $g(\theta, \varphi_0)$ defined by (4.15) is not uniformly monotone, but just monotone, i.e., formally $\gamma_g = 0$, using a projection P on the orthogonal complement of the null space of $\nabla^2 j(\theta^*)$ might facilitate the*

4.3. Local stability and convergence to a local optimum

proof of convergence on this subspace. A possible Lyapunov function in this case is given by $\mathcal{V}(r) = \frac{1}{2}\|Pg(\theta, \varphi_0(\theta))\|^2 + \frac{1}{2}\|u_0 - \varphi_0(\theta)\|_X^2$. Denoting the smallest positive eigenvalue of $\frac{dg}{d\theta}(\theta, \varphi_0(\theta))$ with $\underline{\mu}$, we require that $\xi^T P \frac{dg}{d\theta}(\theta, \varphi_0(\theta)) \xi \leq -\underline{\mu} \|P\xi\|^2$ for all $\xi \in \mathbb{R}^{n_\theta}$. Then a retraction factor λ should be chosen $\lambda > \lambda^* = \frac{\hat{L}}{\gamma_c} + \frac{\hat{L}_g^2}{4\gamma_c \underline{\mu}}$, with

$$\hat{L} = \begin{cases} \frac{((S_0(\theta, u_0) - S_0(\theta, \varphi_0(\theta)))g(\theta, u_0), u_0 - \varphi_0(\theta))_{V^*, V}}{\|u_0 - \varphi_0(\theta)\|_V^2}, & \text{if } u_0 \neq \varphi_0(\theta) \\ 0, & \text{else,} \end{cases} \quad (4.25)$$

and

$$\hat{L}_g = \begin{cases} \frac{g(\theta, \varphi_0(\theta))^T (g(\theta, \varphi_0(\theta)) - g(\theta, u_0))}{\|u_0 - \varphi_0(\theta)\|_V \|g(\theta, \varphi_0(\theta))\|}, & \text{if } u_0 \neq \varphi_0(\theta) \text{ and } g(\theta, \varphi_0(\theta)) \neq 0 \\ 0, & \text{else.} \end{cases} \quad (4.26)$$

However, the null space of $\nabla^2 j(\theta^*)$ depends on the unknown optimal parameter θ^* and can in general not be assessed a priori, thus leaving this approach for further investigation. Alternatively, one can still use regularization (Engl et al., 2000), e.g., by adding a term $\alpha(\theta - \theta_p)^T \Gamma^{-1}(\theta - \theta_p)$ with positive definite Γ and positive α to the cost function $\mathcal{J}(\theta, u_0)$ in (4.5) for which (4.17) yields a minimizer $(\theta^*(\alpha), u_0^*(\alpha))$. Regularization theory provides convergence of $\theta^*(\alpha)$ to a parameter $\bar{\theta}$ that is consistent with the observations as $\alpha \rightarrow 0$ (Engl et al., 2000).

4.3.2. Elliptic and parabolic PDE constraints

As we consider the partially reduced form of the optimization problem with elliptic and parabolic PDE constraints (4.18), the results established for the elliptic problem can be easily transferred given the existence of a solution operator for the parabolic problem.

Theorem 4.3.2. *Let Assumptions 2.1.1, 4.3.1-4.3.6 be satisfied with g replaced by \tilde{g} according to (4.19). Then there exists a $\lambda^* > 0$ such that for all $\lambda > \lambda^*$ solutions to*

$$\begin{aligned} \frac{d\theta}{dr}(r) &= \tilde{g}(\theta, u_0), & \theta(0) &= \theta_0 \\ \frac{du_0}{dr}(r) &= S_0(\theta, u_0)\tilde{g}(\theta, u_0) + \lambda C_0(\theta, u_0), & u_0(r) &= u_{0,0} \end{aligned} \quad (4.27)$$

are well-defined for all $r > 0$ and the local minimizer (θ^*, u_0^*) of the optimization problem (4.4) is a locally exponentially stable steady state of the system (4.27).

With the setting introduced in the last paragraph of Section 4.2, the result directly follows from Theorem 4.3.1.

4.4. Application

To illustrate the continuous analogue of the descent method, we apply it to study CC-chemokine ligand 21 (CCL21) gradient formation in biological tissues. This process is highly relevant in immune responses (Alvarez et al., 2008; Mellman and Steinman, 2001) and described by a reaction-diffusion equation (Hock et al., 2013). In the following, we outline the model, estimate its parameters using the approach proposed in this chapter and interpret the results.

4.4.1. Model formulation

CCL21 gradients are necessary for the guidance of dendritic cells towards lymphatic vessels (Schumann et al., 2010). They are formed by the combination of several biological processes. The chemokine CCL21 is produced in the lymphatic vessels, which cover a subset domain Ω_L of the domain Ω of interest, $\Omega_L \subset \Omega$. The source term is defined via the function

$$Q(x) = \begin{cases} 1, & \text{for } x \in \Omega_L \\ 0, & \text{otherwise.} \end{cases}$$

The concentration of free CCL21 is denoted by u . Free CCL21 binds to a sugar whose concentration is denoted by s . The binding yields immobilized CCL21 whose concentration is denoted by c . The parameters k_1 , k_{-1} , D , γ , and α denote the binding and unbinding rates, the diffusion coefficient, the degradation rate, and the production rate of CCL21 from the lymphatic vessels, respectively. A PDE model for the process has been developed in (Hock et al., 2013) and is given by

$$\begin{aligned} u_t - D\Delta u &= \alpha Q - k_1 us + k_{-1}c - \gamma u \\ \dot{s} &= -k_1 us + k_{-1}c \\ \dot{c} &= k_1 us - k_{-1}c \end{aligned} \tag{4.28}$$

for $t \in]0, T[$ and $x \in \Omega$, with initial conditions $u(0, x) = c(0, x) = 0$, $s(0, x) = s_0$ and no-flux boundary conditions $\frac{\partial}{\partial \nu} u = 0$ where ν is the outer normal on Ω . The parameter s_0 denotes the initial sugar concentration.

4.4. Application

As the formation of the gradient is fast, we consider the steady state of (4.28). With $s + c = s_0$, the steady state for s and c is given by $c = \frac{s_0 u_0}{1+u_0}$ and $s = \frac{s_0}{u_0+1}$ where $u_0 := \frac{k_1 u}{k_{-1}}$ denotes the scaled CCL21 concentration. Using the additional reformulation $\tilde{D} = \frac{D}{\gamma}$, $\tilde{\alpha} = \frac{\alpha k_1}{\gamma k_{-1}}$, the scaled steady-state concentration of CCL21, u_0 , has to fulfill $0 = D\tilde{\Delta}u_0 + \tilde{\alpha}Q - u_0$ and the boundary conditions $\frac{\partial}{\partial\nu}u_0 = 0$.

For the considered process, imaging data have been collected (Weber et al., 2013). These images provide information about the localization of the lymphatic vessels (encoded in Q) and the concentration of immobilized CCL21. As the measured intensity values are corrupted by background fluorescence and as the data are not normalized, we model the readout following (Hock et al., 2013) as

$$y_i = s_l \left(b + \int_{A_i} c(t, x) dx \right),$$

where b denotes the intensity of the background fluorescence, s_l is a scaling constant and $A_i \subset \Omega$ is the domain of the pixel k . As the parameters are structurally non-identifiable, we reformulate the models in terms of $\tilde{b} = s_l b$ and $\tilde{s}_0 = s_l s_0$ in the parameter estimation to one parameter and just consider \tilde{b} and \tilde{s}_0 .

The optimization problem is then given by

$$\begin{aligned} \min_{\theta, u_0} J(\theta, u_0) &= \frac{1}{2} \left\{ \sum_{i=1}^M \log(2\pi\sigma_i^2 \bar{y}_i^2) + \left(\frac{\log(\bar{y}_i) - \log(y_i)}{\sigma_i} \right)^2 \right\} \\ \text{s.t.} \quad &\begin{cases} -\tilde{D}\Delta u_0 + u_0 = \tilde{\alpha}Q, & x \in \Omega \\ \frac{\partial}{\partial\nu}u_0 = 0, & x \in \partial\Omega \\ y_i = \tilde{b} + \int_{A_i} \frac{\tilde{s}_0 u_0(x)}{u_0(x)+1} dx = \tilde{b} + \tilde{s}_0 h_i(u_0) \end{cases} \end{aligned} \quad (4.29)$$

where $\Omega \subseteq \mathbb{R}^2$, σ_i is the scale parameter of the log-normally distributed measurement error and $h_i(u_0) = \int_{A_i} \frac{u_0(x)}{u_0(x)+1} dx$, $i = 1, \dots, M$. The parameter vector θ is given by $\theta = (\tilde{D}, \tilde{\alpha}, \tilde{s}_0, \tilde{b}, \sigma) \in \mathbb{R}^{n_\theta}$, with $n_\theta = 5$.

All parameters are assumed to be non-negative due to their biological meaning. The spaces V and V^* for which we examine the problem are $V = H^1(\Omega)$ and $V^* = H^1(\Omega)^*$. The operator C_0 is given by $C_0(\theta, u_0) = \tilde{D}\Delta u_0 - u_0 + \tilde{\alpha}$. For these spaces and operators it can be checked that all assumptions for applying the method (4.17) are satisfied.

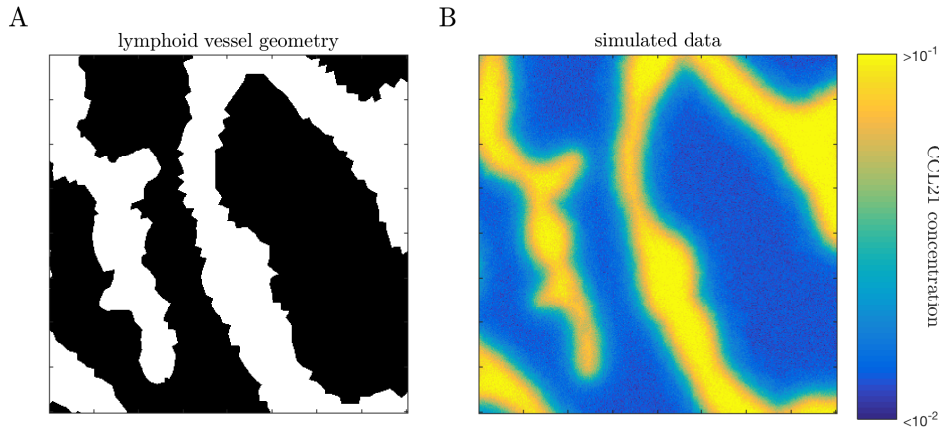


Figure 4.3.: (A) Geometry of a lymphoid vessel obtained from biological imaging data (Weber et al., 2013). (B) Simulated data of the CCL21 gradient generated by simulating model (4.28).

4.4.2. Numerical implementation

For the numerical simulation of the biological process, we employed a finite element discretization of the PDE model. The discretization was obtained using the MATLAB PDE toolbox and accounts for the topology of the model (Figure 4.3 A). The mesh consists of 2170 elements and the concentrations in these elements are the state variables of the discretization. For parameter optimization using the coupled ODE-PDE model (4.17), the same mesh was employed and the states of the discretized PDE were coupled with the ODE for the parameters. This yields a model with $2170 + 5$ equations. The continuous analogue method for parameter estimation was implemented in MATLAB extending the routine published in (Fiedler et al., 2016). The numerical simulation was performed using the MATLAB ODE solver `ode15s`, an implicit scheme applicable to stiff problems. To accelerate the calculations, we implemented the Jacobian of the coupled ODE-PDE model. The simulation of the continuous analogue was terminated, if the gradient of the right-hand side became small, i.e., $\|C_0(\theta, u_0)\|_{V^*} / \|u_0\|_V \leq 10^{-6}$. Furthermore, simulations were interrupted whenever the objective function value became complex, which can happen due to the log-transformation of the output.

4.4.3. Simulated data

To evaluate the convergence properties of the proposed algorithm for this model, we considered published simulated data for the ground truth (Hross, 2016; Hross et al., 2018).

Table 4.1.: True parameters, estimated parameters and parameter ranges for the latin hypercube sampling for the CCL21 model.

name	true value	estimates	lower bound sampling	upper bound sampling
D	$8.50 \cdot 10^1$	$8.51 \cdot 10^1$	$1.50 \cdot 10^{-2}$	$2.50 \cdot 10^4$
α	$2.40 \cdot 10^{-1}$	$2.36 \cdot 10^{-1}$	$4.50 \cdot 10^{-5}$	$2.00 \cdot 10^1$
s_0	$2.10 \cdot 10^{-1}$	$2.13 \cdot 10^{-1}$	$2.50 \cdot 10^{-3}$	$3.00 \cdot 10^0$
b	$1.30 \cdot 10^{-2}$	$1.30 \cdot 10^{-2}$	$2.50 \cdot 10^{-3}$	$1.00 \cdot 10^0$
σ	$5.00 \cdot 10^{-2}$	$4.99 \cdot 10^{-2}$	$2.50 \cdot 10^{-3}$	$3.00 \cdot 10^0$

The geometry of lymphatic tissue was extracted from the available imaging data (Weber et al., 2013) using the MATLAB PDE toolbox. On this geometry the discretized PDE was simulated using biologically plausible parameter values (Table 4.1). The simulated data for CCL21 gradient formation process were corrupted by noise to obtain a plausible scenario (Figure 4.3 B).

4.4.4. Optimization

The objective function for most parameter estimation problems is non-convex and can be multi-modal. For this reason, we employed multi-start local optimization using the continuous analogue for which we have established local convergence in this chapter. The starting points for the local optimizations were sampled using a latin hypercube approach with lower and upper bounds provided in Table 4.1. We used a linear parametrization for the states and a log-parametrization $\xi = \log(\theta)$ for the parameters following previous evaluations for biochemical systems (Raue et al., 2013b). We did not implement any bounds for values of parameter or states. The implementation of the multi-start local optimization is based upon the MATLAB toolbox PESTO (Stapor et al., 2018b). The implementation of the objective function and finite element schemes were adapted from (Hock et al., 2013). For the local optimization with the continuous analogue we chose the negative gradient as descent direction.

As a reference, we performed also multi-start local optimization using a discrete iterative optimization method. We used the state-of-the-art optimizer `fmincon.m` with the starting points sampled for the continuous analogue and the interior point algorithm implemented in the MATLAB Optimization Toolbox. This interior point algorithm employs either a Newton step, where the Hessian is approximated by the Broyden-Fletcher-Goldfarb-Shanno (BFGS) algorithm, or a conjugate gradient step using a trust region (Byrd et al.,

1999, 2000; Waltz et al., 2006). The optimizer was provided with the objective function, the nonlinear constraint, as well as the corresponding derivatives. We used the same parametrization as for the continuous analogue and additionally constrained parameter values for $\xi = \log(\theta)$ in the optimization by the same upper and lower bounds used for the sampling (Table 4.1). The value of u_0 at the nodes of the mesh for the finite element discretization was constrained using upper and lower bounds for the optimization to lie in $[-1, 3]$. A total of 2000 iterations and 4000 function evaluations was allowed.

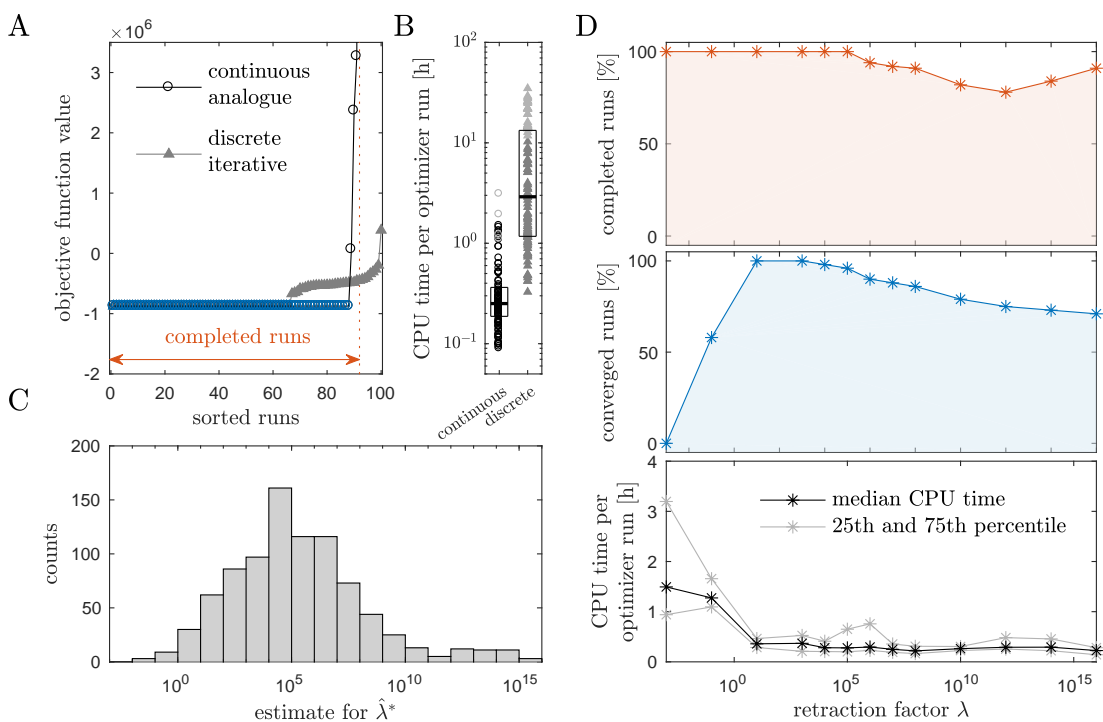


Figure 4.4.: Results of parameter estimation for CCL21 model. (A) Sorted objective function values for the multi-start optimization with continuous analogue ($\lambda = 10^7$) and discrete iterative procedure. Converged runs are indicated in blue. (B) CPU time needed per optimizer run for the optimization using the continuous analogue and the discrete iterative procedure (lighter grey color indicates runs which stopped because the maximal number of iterations was reached). The box covers the range between the 25th and the 75th percentile of the distribution. The median CPU time is indicated by a line. (C) Histogram of values for $\hat{\lambda}^*$ (Remark 4.3.2) obtained for 1000 points sampled in parameter-state space. (D) Percentage of completed runs (top), converged runs (middle), and median as well as 25th and 75th percentile of the runtime of completed runs (bottom) for different values of λ . For each value of λ , 100 local optimization runs were performed.

4.4.5. Comparison of continuous analogue and discrete iterative procedure

We performed 100 local optimization runs with the continuous analogue for the retraction factor $\lambda = 10^7$ and a discrete iterative method (Figure 4.4 A). Both methods found the same best parameter value (Table 4.1) and achieved a good fit to the data. The assessment of the results revealed a good convergence of the continuous analogue. Almost 90% of the runs achieved an objective function value which was comparable with the best objective function value found across all runs (relative difference $< 0.001\%$). Overall, 96% of the runs finished successfully, meaning that either the optimization was stopped because the stopping criterion was fulfilled or the maximum number of iterations was reached, while 4% of the runs stopped prematurely.

The discrete iterative optimization converged for 66% of the runs to the optimal value (Figure 4.4 A). Accordingly, the success rate was substantially lower than for the proposed continuous analogue. Of the runs which did not converge to the global optimum 25 runs were stopped because the maximal number of iterations was reached.

For the considered problem, the continuous analogue outperformed the discrete iterative method regarding the CPU time (Figure 4.4 B). We found a median CPU time of 15 minutes for the continuous solver and 174 minutes for the discrete iterative procedure. In light of the fact that the discrete iterative method uses second order information, it is interesting to observe that a continuous analogue using the negative gradient is more efficient. One possible explanation is that the efficiency of the continuous analogue is a result of the application of sophisticated numerical solvers. The adaptive, implicit solver `ode15s`, which is provided with the analytical Jacobian of the ODE-PDE model, might facilitate large step-sizes and fast convergence. Indeed, the Jacobian also provides second order information.

4.4.6. Evaluation of retraction factor influence

As an analytical calculation of the bound for the retraction factor was not possible, we sampled 1000 points in parameter-state space and evaluated the estimate for the lower bound $\hat{\lambda}^*$ (Remark 4.3.2) to address Problem 4.1.2. The histogram of the resulting values for $\hat{\lambda}^*$ is presented in Figure 4.4 C. The values for $\hat{\lambda}^*$ span many orders of magnitude, and the distribution peaks at 10^4 . This result indicated that for different starting points very different retraction factors might be ideal.

To investigate the convergence properties for the different values of the retraction factor λ , we performed 100 local optimization runs for a range of different retraction factors. For each retraction factor we assessed the number of completed runs and the number of converged runs (Figure 4.4 D). Interestingly, as λ was increased the percentage of completed runs decreased. Yet, for large retraction factors many of the completed runs also converged, while for small retraction factors no runs converged as the maximum number of iterations becomes too large. The median CPU time for the optimization of one run decreased for increasing values of λ (Figure 4.4 D). Notably, for the small values, the median CPU time was nearly six to seven times higher than the smallest one. The quantiles indicate that also the variability was higher for small values of λ . These results indicated that the retraction factor should be chosen large enough but not too large.

In summary, the analysis of the model of CCL21 gradient formation revealed that the retraction factor λ has a substantial influence on the convergence properties as well as the run time. For low values of λ starts did not converge while for large values of λ increasing stiffness of the problem could be observed. In an intermediate regime, which could here also be found by random sampling, we found the best convergence properties.

4.5. Discussion

Parameter estimation for PDE models is an important problem in a wide range of applications. Robustness and performance of the available iterative methods is however often limited. Hence, we investigated a continuous analogue approach. To address Problem 4.1.1, we introduced continuous analogues of descent methods for optimization with PDE constraints and proved local convergence of their solutions to the optima.

We demonstrated the applicability of continuous analogues for a model of gradient formation in biological tissues and compared them with an iterative discrete procedure. The results highlight the potential of the continuous analogues, e.g., a high convergence rate and lower computation times than the discrete iterative procedure. For the comparison we used the MATLAB optimization routine `fmincon.m`, a state-of-the-art discrete iterative procedure. Alternatives would be IPOPT or KNITRO. As `fmincon.m` is a generic interior point method, there might be approaches which are efficient for the considered PDE constrained problems (see also the no free lunch theorem (Wolpert and Macready, 1997)). To address the second research problem (Problem 4.1.2) we evaluated the influence of the retraction factor and revealed the importance of an appropriate choice of the retraction factor as well as the issue of premature stopping. We provided a lower bound for λ which

ensures local convergence. As this bound might however be conservative and can only be assessed point-wise, the use of adaptive methods might be interesting. To address the issue of premature stopping, bounds for parameters and state variables have to be implemented, e.g., by including log-barrier functions (Boyd and Vandenberghe, 2004) in the objective function or through projection into the feasible space.

In the application problem we only considered elliptic PDE constraints as for the proposed continuous analogues parabolic constraints can be encapsulated in the objective function. This changes the objective function landscape and indirectly influences the convergence. Conceptually, it should also be possible to formulate continuous analogues which do not require a solution operator for the parabolic PDE but also have the solution of the parabolic PDE as a state variable. This mathematically more elegant approach is left for future research.

In conclusion, in this chapter we presented continuous analogues for a new problem class. Similar to other problem classes for which continuous analogues have been established (Fiedler et al., 2016; Tanabe, 1985), we expect an improvement of convergence and computation time.

The method and its analysis apply as they are to the case of infinite dimensional parameters θ . However, in that situation, the inverse problem of identifying θ is often ill-posed, so the assumption of practical identifiability (cf. Assumption 4.3.5 and Remark 4.3.3) might not be satisfied. To restore stability, regularization can be employed.

Chapter 5.

Modeling of population dynamics on reduced data spaces

While we considered reduced parameter spaces in the previous chapters, in this chapter we consider mathematical models using reduced data spaces. We investigate the dynamics of heterogeneous cell populations using observations in the high-dimensional data space generated by single-cell experiments. To facilitate the modeling of the population dynamics, we employ a dimension reduction of the high-dimensional data to one-dimensional cell state trajectories. Firstly, we introduce the biological background of the process we want to model and the structure of the data that experiments typically provide (Section 5.1). Then, we develop the mathematical model to describe the processes in the reduced data space, a diffusion-advection-reaction equation, and discuss the numerical implementation (Section 5.2). In Section 5.3, we describe the model parameterization and introduce an objective function for the parameter estimation. We illustrate the mathematical model and the parameter estimation framework using T-cell maturation data (Section 5.4).

This chapter is based on and partly identical to the manuscript *Inferring population dynamics from single-cell RNA-sequencing time series data* by David S. Fischer, Anna K. Fiedler, Eric M. Kernfeld, Ryan M.J. Genga, Aimée Bastidas-Ponce, Mostafa Bakhti, Heiko Lickert, Jan Hasenauer, Rene Maehr and Fabian J. Theis published in *Nature Biotechnology*, 37(4):461–468, 2019 (Fischer et al., 2019).

5.1. Introduction and problem formulation

In this section, we briefly describe the biological background of cellular development and how single-cell experiments, especially using single-cell messenger RNA (mRNA) sequenc-

ing, are used to study this process. We sketch how the sequencing data is processed to reduce its dimension and explain the structure of the data we consider in the remainder of the chapter. Based on this information we conclude the section by formulating the research problem. A schematic representation of the whole analysis process including dimension reduction, dynamical modeling, parameter estimation and hypothesis testing is depicted in Figure 5.1.

5.1.1. Biological background

In this chapter, we study population dynamics consisting of differentiation and proliferation, i.e., cellular development, for heterogeneous cell-populations. During differentiation processes a cell population moves through different stages, e.g., cells go from unspecialized stem or progenitor cells to more specialized ones while also dividing or possibly dying (Figure 5.2 A). The described processes of cellular development occur for example prenatally in embryogenesis but also in adult hematopoiesis. Often, a progenitor cell can give rise to several specialized cell types. Previously, cell types were commonly assigned to cells by measuring the expression of surface markers or marker genes (e.g. Yui and Rothenberg (2014)). However, technical advances permit the measurement of high-dimensional cellular features on a single-cell level and a characterization of cell types using these features. In the following, we will focus on single-cell sequencing as an example of those experimental techniques. In recent years, single-cell sequencing technologies have been established that enable gene expression measurements from a mixture of cells with single-cell resolution. This means we can measure the transcriptome, i.e., the number of mRNA molecules transcribed from the genome, in single cells for hundreds to thousands of cells at the same time (Klein et al., 2015). This technology permits to look at development as a continuous trajectory in transcriptome space (Figure 5.2 B), i.e., the space where each gene transcript is a dimension and a cell's position is assigned by the amount of transcripts of the genes it contains. As cell development is usually asynchronous, cells in different differentiation stages can be found in a tissue sample at the same measurement time point (Figure 5.2 C). Combining several snap shots of different time points comprises more information about the differentiation process in transcriptome space. As the transcriptome space is very high dimensional (10^4 to 10^6 dimensions) (Klein et al., 2015), it is helpful to reduce the dimension by extracting the main developmental trajectory using pseudo temporal orderings and assigning a continuous cell state to the cells (Figure 5.2 D). This chapter considers diffusion maps (see, e.g., Haghverdi et al. (2015)) and diffusion pseudotime (Haghverdi et al., 2016) for dimension reduction. If the development can potentially give rise to more than one terminal cell fate, the trajectory branches and cells are assigned to a branch. Using

5.1. Introduction and problem formulation

correlation of the diffusion pseudotime along trajectories, branching can be identified in the data (Haghverdi et al., 2016). As the diffusion pseudotime is based on an ordering of the measured single cells, it can be scaled without changing the biological interpretation. In the following, we use diffusion pseudotime as the cell state, however this method is not limited to this particular pseudotime.

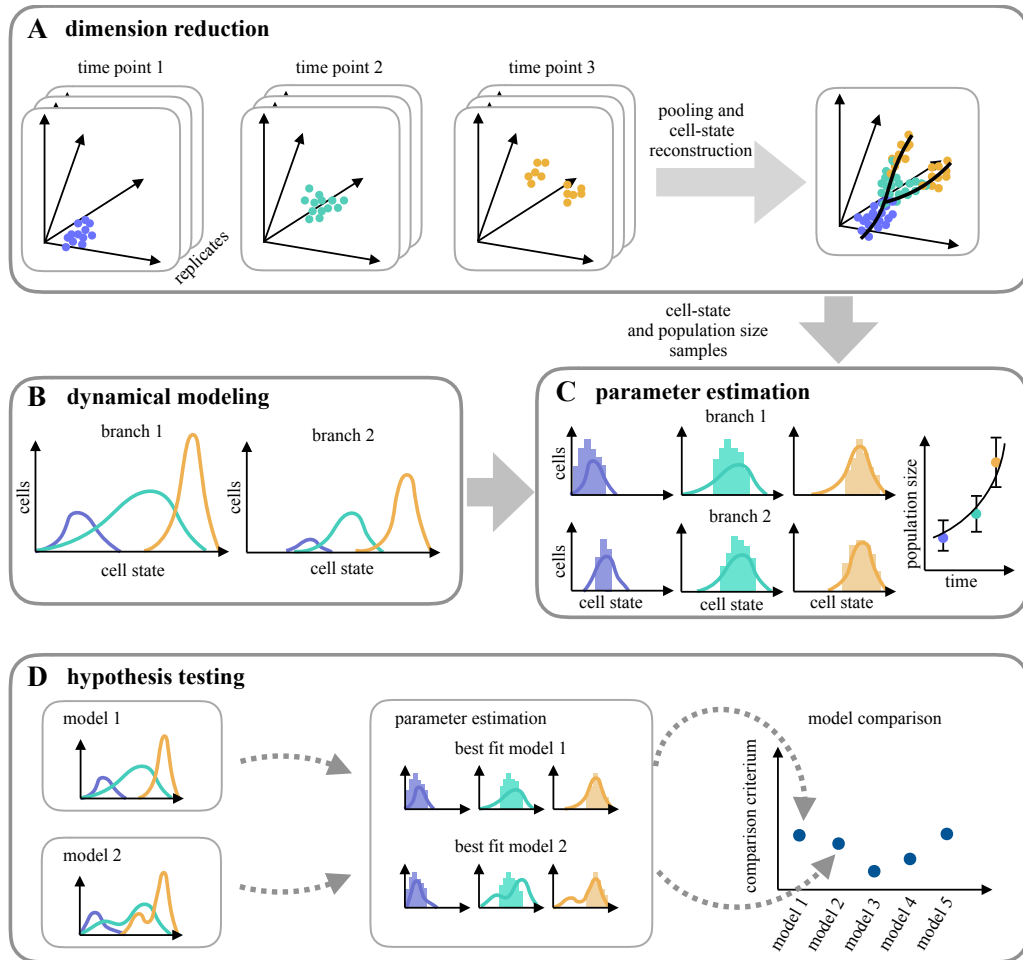


Figure 5.1.: Schematic illustration of analysis pipeline for single-cell experiments using reduced data spaces. (A) The data dimension is reduced by pooling the individual single-cell experiment replicates at different time points, reconstructing cell state trajectories using the pooled sample and assigning cell states along the trajectory to the individual cells. (B) The population dynamics in the reduced data space are described using dynamical models of the cell distribution along the cell state trajectory. (C) The dynamical model is calibrated to the data by estimating the model parameters that optimize the fit of data and model. (D) Competing hypotheses are tested by estimating the model parameters of the competing models and using model comparison criteria to select the models with substantial support.

5.1.2. Data structure

The measurements typically consist of single-cell sequencing snapshot data and population size snapshot data, i.e., the data represents the state of the system at particular time points but does not capture individual trajectories. At a number of single-cell measurement time points $t \in T^{cdf}$, where T^{cdf} denotes the set of measurement time points, a subset of cells of the total population is sequenced. The single-cell sequencing measurements at time point t are replicated $|R_t|$ times using independent cell populations, where R_t denotes the set of replicates at time point t . The number of replicates might differ between time points. The cell state $s \in [0, 1]$ is computed from the mRNA content of each cell, e.g., using diffusion pseudo time. Further, branching is identified and if there is branching a branch membership out of the set of branches, $b \in B$, is assigned to each cell. This procedure yields a set of cell state measurements $S_{b,t}^r$ for a replicate $r \in R_t$ at time point t and on branch b . If the process branches, the weight of branch b , i.e., the fraction of cells on branch b compared to the total amount of measured cells, $w_{b,t}^r = \frac{|S_{b,t}^r|}{\sum_{b \in B} |S_{b,t}^r|}$, at time t is calculated. The population size data is obtained from different individuals than the cell state measurements. Therefore, it is possible to have measurements at a different set of measurement time points $t \in T^N$ and a different set of replicates Z_t . The number of replicates $|Z_t|$ can differ from time point to time point. The population size measurement for the replicate z at time point t is denoted by N_t^z .

5.1.3. Problem formulation

Previous models of cell differentiation and development were based on discrete cell types and were therefore modeled using compartment models. Discrete cell types are assigned to a cell using gating (Figure 5.2 E), i.e., prior knowledge of the biological process is used to define regions in some feature space, e.g., surface marker concentration, fluorescence or transcriptomic expression, that correspond to a cell type. The amount of cells in each region is counted and used as observable. In this chapter, we want to use the continuous cell state measurements to elucidate the population dynamics during development and combine trajectory and real time information (Figure 5.2 F).

Problem 5.1.1. *Develop a mathematical model to describe cell proliferation and differentiation using a reduced data space.*

The combination of different data types in one likelihood is usually not straight forward, especially, if the data types differ in information content and reproducibility. The inherent

5.1. Introduction and problem formulation

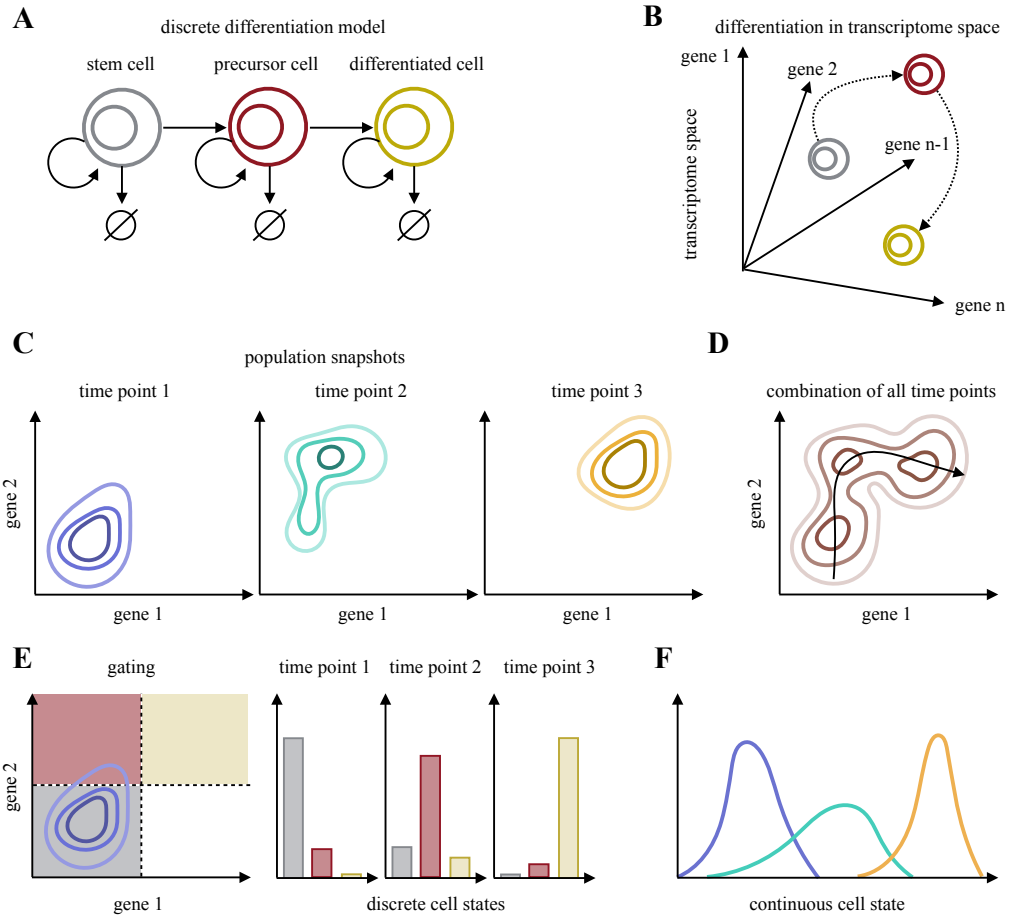


Figure 5.2.: Illustration of development in transcriptome space and discrete and continuous cell state models (A) Discrete view on development using cell state compartments. (B) Differentiation as trajectory in the high dimensional transcriptome space. (C) Schematic illustration of snapshots of the cell state distribution of a population in a two dimensional transcriptome space for three time points. (D) Distribution of the combination of all snapshots together with one-dimensional trajectory. This trajectory is equated with the continuous cell state. (E) Illustration of gating to assign discrete cell states for the population distribution. Gates define regions in the transcriptome space that correspond to a certain cell type. For each time point the amount of cells in each state is determined and used as data. (F) Distribution of cells in the population along the continuous cell state for the three measurement time points.

differences between population level measurements, like population size and single-cell measurements, e.g., cell state, have to be taken into account. Further, a suitable error model for the single-cell measurements has to be developed.

Problem 5.1.2. *Develop a likelihood for parameter estimation using single-cell sequencing measurements and total population size.*

5.2. Mathematical model

In this section, we address Problem 5.1.1 and develop models that describe the proliferation and differentiation dynamics of a population of cells by using the reduction to a one dimensional cell state from the high dimensional single-cell RNA sequencing data. A partial differential equation in time and cell state is used to describe the dynamics of the cell state distribution of single cells in the population. We start by considering a differentiation process that only terminates in a single fate. Afterwards, we extend the model to differentiation processes with branching.

5.2.1. Non-branching model

To develop the non-branching model, we consider the bounded cell state interval $\Omega \subset \mathbb{R}$ as the state space and we choose the time interval $I =]0, T[$ for some maximal duration $0 < T < \infty$. Further, we model the dynamics of the number density, $u(s, t): \Omega \times I \rightarrow \mathbb{R}$ of cells in a certain cell state $s \in \Omega$ at time $t \in I$. The integral $\int_{s_1}^{s_2} u(s, t) ds$ of the number density over a cell state interval $[s_1, s_2]$ provides the number of cells in $[s_1, s_2]$ at time point t . Assuming directed and random movement of cells in cell state as well as cell division and cell death, the temporal evolution of $u(s, t)$ can be described using population balance as a PDE. Differentiation is modeled as a deterministic process that is represented by a drift term with parameter $v(s, t) \in \mathbb{R}^+$. We further include a diffusion term with parameter $D(s, t) \in \mathbb{R}^+$ that describes cell intrinsic stochasticity in RNA transcription as well as fluctuations in cell state that might arise through the projection of cells on the one dimensional cell state. The proliferation dynamics, i.e., cell division and cell death, are included in a functional term with proliferation rate $g(s, t) \in \mathbb{R}$ that might act as a source or a sink depending on whether cell division or cell death dominates the proliferation dynamics. It is further assumed that cell division does not alter the cell state, meaning that the daughter cells possess at birth the same cell state as the mother cell. All of the above rates depend on the cell state and the time. The combination of the outlined mechanisms leads to a reaction diffusion advection equation,

$$\frac{\partial}{\partial t} u(s, t) = \underbrace{\frac{\partial}{\partial s} (D(s, t) \frac{\partial}{\partial s} u(s, t))}_{\text{diffusion}} - \underbrace{\frac{\partial}{\partial s} (v(s, t) u(s, t))}_{\text{drift}} + \underbrace{g(s, t) u(s, t)}_{\text{population growth}} \quad \text{in } \Omega \times I. \quad (5.1)$$

As the cell state can be rescaled, we will consider a cell state interval $\Omega =]0, 1[$.

To assess the proliferation dynamics, we also have to consider the total population size, $N(t)$, that is calculated as the integral of $u(s, t)$ with respect to the cell state s over the whole state domain $[0, 1]$,

$$N(t) = \int_0^1 u(s, t) ds. \quad (5.2)$$

The boundary conditions generally depend on the biological process and the experimental setting in consideration. Throughout this chapter, we assume no flux boundaries, i.e., no cells enter or leave the studied cell state interval through the boundaries,

$$\left(D(s, t) \frac{\partial u}{\partial s}(s, t) - v(s, t)u(s, t) \right) \Big|_{s=0} = 0 \text{ and } \frac{\partial u}{\partial s}(s, t) \Big|_{s=1} = 0 \quad \forall t \in [0, T]. \quad (5.3)$$

At the right boundary, $s = 1$, we assume that the cells are differentiated and do not advance further in the cell state leading to zero drift, $v(1, t) = 0$ and use the corresponding no flux boundary condition.

The initial condition depends on the experimental setup and the biological process under consideration. In this chapter, we consider initial conditions, $u(s, 0) = u_0(s)$, that are derived from the data at $t = 0$ directly by approximating the initial density from the data using kernel smoothing with a normally distributed kernel and scaling the resulting density with the initial population size. To discuss the most general case, we describe the steps we carried out for an example with repeated measurements at $t = 0$. First, the kernel density estimate (kde) of the set of cell state measurements S_0^r is computed on a fixed grid for each replicate $r \in R_0$. Then we normalized the kde yielding an initial probability distribution denoted by $\text{kde}(s|S_0^r)$ and computed the point-wise mean of the normalized kdes on the grid,

$$\frac{1}{|R_0|} \sum_{r \in R_0} (\text{kde}(s|S_0^r)). \quad (5.4)$$

The initial number density is given as

$$u_0(s) = \frac{1}{|R_0|} \sum_{r \in R_0} (\text{kde}(s|S_0^r)) \bar{N}_0, \quad (5.5)$$

with $\bar{N}_0 = \frac{1}{|Z_0|} \sum_{z \in Z_0} N_0^z$ denoting the mean of the population size measurement replicates $z \in Z_0$ at time point zero. Alternatively, one could also use the normalized kde of the pooled sample, $S_0 := \bigcup_{r \in R_0} S_0^r$ scaled by \bar{N}_0 . This particular computation of the initial condition from the data requires population size and single-cell measurements at $t = 0$. Apart from estimating the initial distribution from the data, one could also add a parameter dependent mixture of normal distributions and estimate these parameters together

with the other unknowns of the model. This approach is useful if there are no single-cell measurements at $t = 0$ or the assumption is biologically plausible.

5.2.2. Branching model

To extend the model to a differentiation process with two terminal cell fates, we consider coupled PDEs on two branches, a main branch and a side branch, that overlap in a branching region. The main branch, comprises the whole interval $\Omega_1 =]0, 1[$, while the side branch starts at the beginning of the branching region, $s_a \in]0, 1[$ and consists of the interval $\Omega_2 =]s_a, 1[$ (Figure 5.3 A). In the branching region cells can switch from branch i to the other branch j with propensities δ_{ji} . The number density of cells on branch $b \in \{1, 2\}$ is denoted by $u_b(s_b, t): \Omega_b \times I \rightarrow \mathbb{R}$. The dynamics on each branch consist of the same processes as in the non-branching case and additionally the switch of cells between branches in the branching region. The individual rates might differ between branches indicated by a subscript $b \in \{1, 2\}$. The resulting PDE is given by,

$$\begin{aligned} \frac{\partial}{\partial t} u_1(s_1, t) &= \frac{\partial}{\partial s_1} (D_1(s_1, t) \frac{\partial}{\partial s_1} u_1(s_1, t)) - \frac{\partial}{\partial s_1} (v_1(s_1, t) u_1(s_1, t)) + g_1(s_1, t) u_1(s_1, t) \\ &\quad - \mathbb{1}_{[s_a, s_e]}(s_1) (\delta_{12} u_1(s_1, t) - \delta_{21} u_2(s_1, t)) \text{ in } \Omega_1 \times I \\ \frac{\partial}{\partial t} u_2(s_2, t) &= \frac{\partial}{\partial s_2} (D_2(s_2, t) \frac{\partial}{\partial s_2} u_2(s_2, t)) - \frac{\partial}{\partial s_2} (v_2(s_2, t) u_2(s_2, t)) + g_2(s_2, t) u_2(s_2, t) \\ &\quad + \mathbb{1}_{[s_a, s_e]}(s_2) (\delta_{12} u_1(s_2, t) - \delta_{21} u_2(s_2, t)) \text{ in } \Omega_2 \times I \end{aligned} \quad (5.6)$$

in which the indicator function $\mathbb{1}_{[s_a, s_e]}(s_b)$ defines the branching interval $[s_a, s_e] \subset]0, 1[$,

$$\mathbb{1}_{[s_a, s_e]}(s_b) = \begin{cases} 1, & \text{if } s_b \in [s_a, s_e] \\ 0, & \text{else,} \end{cases} \quad (5.7)$$

where s_e denotes the end of the branching region. Depending on the number and the geometry of branches that is determined in the preprocessing step, this model can in general be extended to a set of branches B with $\|B\| > 2$.

The boundary conditions for two branches are chosen as non-flux boundaries,

$$\begin{aligned} \left(D_1(s_1, t) \frac{\partial u_1}{\partial s_1}(s_1, t) - v_1(s_1, t) u_1(s_1, t) \right) \Big|_{s_1=0} &= 0 \text{ and } \frac{\partial}{\partial s_1} u_1(s_1, t) \Big|_{s_1=1} = 0, \forall t \in [0, T] \\ \left(D_2(s_2, t) \frac{\partial u_2}{\partial s_2}(s_2, t) - v_2(s_2, t) u_2(s_2, t) \right) \Big|_{s_2=s_a} &= 0 \text{ and } \frac{\partial}{\partial s_2} u_2(s_2, t) \Big|_{s_2=1} = 0, \forall t \in [0, T]. \end{aligned} \quad (5.8)$$

At the right boundary, $s_b = 1$, we assume again that the cells are differentiated and do not advance further in the cell state leading to zero drift, $v_i(1, t) = 0$.

The initial condition on each branch is computed as in Section 5.2.1 but is additionally scaled with the mean of the fraction of cells measured on the respective branch in the single-cell experiment replicates,

$$\bar{w}_{b,0} = \frac{1}{|R_0|} \sum_{r \in R_0} w_{b,0}^r. \quad (5.9)$$

This yields the initial number density,

$$u_{b,0}(s_b) = \frac{1}{|R_0|} \sum_{r \in R_0} \left(\text{kde}(s_b | S_{b,0}^r) \right) \bar{w}_{b,0} \bar{N}_0. \quad (5.10)$$

The total population size, $N(t)$, is calculated as the sum of the integrals of $u_b(s_b, t)$ with respect to the cell state s_b ,

$$N(t) = \int_0^1 u_1(s_1, t) ds_1 + \int_{s_a}^1 u_2(s_2, t) ds_2. \quad (5.11)$$

5.2.3. Numerical scheme for simulation

To describe the growth dynamics, the spatial discretization of the PDE should conserve mass if no growth dynamics are present $g(s, t) = 0$ and yield the expected dynamics $\dot{N} = g_c N(t)$ if the birth-death rate is given by a constant rate $g(s, t) = g_c$. Hence, we use a finite volume approach for the discretization of the PDE in space (Hundsdoerfer and Verwer, 2003). This discretization together with the method of lines (Hundsdoerfer and Verwer, 2003), i.e., considering the ODE resulting from the discretization, is used to simulate the PDE. For the finite volume discretization of the non-branching process (5.1) in one dimension, the cell state interval Ω is divided into n_V control volumes (i.e., intervals) of length h by $n_V + 1$ equally spaced grid points. The nodes located at the center of an interval are denoted with $s_i = (i - \frac{1}{2})h$, $i = 1, \dots, n_V$ and the borders of the interval are $s_{i-\frac{1}{2}} = (i-1)h$ on the left and $s_{i+\frac{1}{2}} = ih$ on the right. For each center node s_i , we consider the average number density in the interval $[s_{i-\frac{1}{2}}, s_{i+\frac{1}{2}}[$

$$u_i(t) = \frac{1}{h} \int_{s_{i-\frac{1}{2}}}^{s_{i+\frac{1}{2}}} u(s, t) ds \quad \text{for } i = 1, \dots, n_V. \quad (5.12)$$

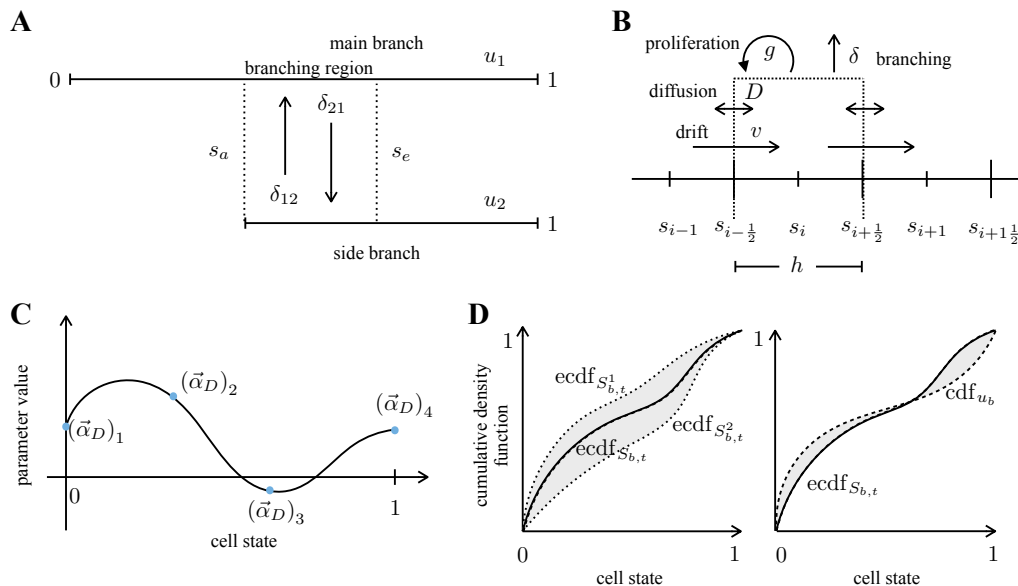


Figure 5.3.: (A) Sketch of the branching model with main branch, side branch and branching region. (B) Illustration of the finite volume discretization including fluxes that contribute to change in the control volume. (C) Visualization of the parametrization using splines. The model parameters are the values at the spline nodes indicated in blue. The rest of the rate is completed using cubic splines between the nodes. (D) Illustrative sketch of the areas between cumulative densities computed for the density fit of the log-likelihoods in Equations 5.31 and 5.33. The ecdf of the pooled sample is shown as a line in both subfigures (black line). The empirical cumulative density functions (ecdfs) of two replicates (dotted lines) and the area between these ecdfs and the ecdf of the pooled sample (grey areas), which are used to compute the statistics, are shown. The cumulative density functions (cdf) of the model simulation is depicted (dashed line) and the area between this cdf and the ecdf of the pooled sample (grey area) is compared to the statistics obtained from the data.

The average number density in an interval changes in time by inflow and outflow at the boundaries of the interval (Figure 5.3 B). For model (5.1) this is approximated by

$$\begin{aligned} \frac{du_i}{dt} = & \frac{1}{h^2} \left(D_{i-\frac{1}{2}} (u_{i-1} - u_i) - D_{i+\frac{1}{2}} (u_i - u_{i+1}) \right) + \\ & \frac{1}{2h} \left(v_{i-\frac{1}{2}} (u_{i-1} + u_i) - v_{i+\frac{1}{2}} (u_i + u_{i+1}) \right) + g_i u_i \text{ for } i \in \{2, n_V - 1\}, \end{aligned} \quad (5.13)$$

for the inner intervals and

$$\frac{du_1}{dt} = -\frac{1}{h^2} D_{1+\frac{1}{2}} (u_1 - u_2) - \frac{1}{2h} v_{1+\frac{1}{2}} (u_1 + u_2) + g_1 u_1 \quad (5.14)$$

at the right boundary $i = 1$ and

$$\frac{du_{n_V}}{dt} = \frac{1}{h^2} D_{n_V-\frac{1}{2}} (u_{n_V-1} - u_{n_V}) - \frac{1}{2h} v_{n_V-\frac{1}{2}} (u_{n_V-1} + u_{n_V}) + g_{n_V} u_{n_V} \quad (5.15)$$

5.2. Mathematical model

at the left boundary $i = n_V$, with $D_{i\pm\frac{1}{2}} = D(s_{i\pm\frac{1}{2}}, t)$, $v_{i\pm\frac{1}{2}} = v(s_{i\pm\frac{1}{2}}, t)$ and $g_i = g(s_i, t)$ (Hundsdofer and Verwer, 2003).

By summing the equations (5.13)-(5.15) and setting $g(s, t) = 0$, one finds

$$\sum_i \frac{du_i}{dt} = \frac{1}{h} \frac{dN}{dt} = 0. \quad (5.16)$$

Hence, this spatial discretization conserves mass, i.e., the populations size stays constant. No mass is created or destroyed by approximating the diffusion and advection dynamics numerically.

For the system with branching (5.6), the finite volume approximation results in a system of two coupled ODEs. We denote the averaged densities in an interval i corresponding to u_1 with $u_{1,i}$ and the ones corresponding to u_2 with $u_{2,i}$. On the side branch, we chose the same interval length h as on the main branch and adapted the number of intervals. If we denote the interval that contains the start of the branching region, s_a , with the n_a , i.e., $s_a \in [s_{n_a-\frac{1}{2}}, s_{n_a+\frac{1}{2}}[$, the side branch contains the intervals indexed by n_a, \dots, n_V and the number of intervals on the side branch is $n_V - n_a + 1$. The intervals in the branching region are indexed with $i \in \{n_a, \dots, n_e\}$ on both branches. As the equations for intervals outside the branching region largely resemble (5.13)-(5.15), we omit them here. The complete ODE system can be found in Appendix B. The temporal evolution of the average number density in the intervals in the branching region on the main branch for the finite volume discretization of (5.6) is given by,

$$\begin{aligned} \frac{du_{1,i}}{dt} = & \frac{1}{h^2} \left(D_{1,i-\frac{1}{2}} (u_{1,i-1} - u_{1,i}) - D_{1,i+\frac{1}{2}} (u_{1,i} - u_{1,i+1}) \right) \\ & + \frac{1}{2h} \left(v_{1,i-\frac{1}{2}} (u_{1,i-1} + u_{1,i}) - v_{1,i+\frac{1}{2}} (u_{1,i} + u_{1,i+1}) \right) \\ & + g_{1,i} u_{1,i} - \delta_{12} u_{1,i} + \delta_{21} u_{2,i} \end{aligned} \quad (5.17)$$

for $i \in \{n_a, \dots, n_e\}$. The temporal evolution of volumes in the branching region at the left boundary $i = n_a$ on the side branch is given by

$$\begin{aligned} \frac{du_{2,n_a}}{dt} = & -\frac{1}{h^2} D_{2,n_a+\frac{1}{2}} (u_{2,n_a} - u_{1,n_a+1}) \\ & - \frac{1}{2h} v_{2,n_a+\frac{1}{2}} (u_{2,n_a} + u_{2,n_a+1}) + g_{2,n_a} u_{2,n_a} + \delta_{12} u_{1,n_a} - \delta_{21} u_{2,n_a} \end{aligned} \quad (5.18)$$

and for $i \in \{n_a + 1, \dots, n_e\}$ by

$$\begin{aligned} \frac{du_{2,i}}{dt} &= \frac{1}{h^2} \left(D_{2,i-\frac{1}{2}} (u_{2,i-1} - u_{2,i}) - D_{2,i+\frac{1}{2}} (u_{2,i} - u_{2,i+1}) \right) \\ &+ \frac{1}{2h} \left(v_{2,i-\frac{1}{2}} (u_{2,i-1} + u_{2,i}) - v_{2,i+\frac{1}{2}} (u_{2,i} + u_{2,i+1}) \right) \\ &+ g_{2,i}u_{2,i} + \delta_{12}u_{1,i} - \delta_{21}u_{2,i}. \end{aligned} \quad (5.19)$$

The initial condition for the ODE is computed from the kernel density approximation $u_{b,0}(s_b)$ in equation 5.10 by using the trapezoidal rule,

$$u_{b,0,i} = 0.5 \left(u_{b,0} \left(s_{i-\frac{1}{2}} \right) + u_{b,0} \left(s_{i+\frac{1}{2}} \right) \right). \quad (5.20)$$

The total population size is computed from the $u_{b,i}$ as

$$N(t) = h \left(\sum_{i=1}^{n_V} u_{1,i}(t) + \sum_{i=n_a}^{n_V} u_{2,i}(t) \right) = N_1(t) + N_2(t), \quad (5.21)$$

where the population size on the main branch is $N_1(t) = h \sum_{i=1}^{n_V} u_{1,i}(t)$ and the population size on the side branch is $N_2(t) = h \sum_{i=n_a}^{n_V} u_{2,i}(t)$. Accordingly, we compute the probability to observe a cell in a certain interval $[s_{i-\frac{1}{2}}, s_{i+\frac{1}{2}}]$ on branch b as,

$$p_{b,i}(t) := \frac{1}{N_b(t)} \int_{s_{i-\frac{1}{2}}}^{s_{i+\frac{1}{2}}} u_b(s_b, t) ds_b = \frac{h}{N_b(t)} u_{b,i}(t). \quad (5.22)$$

5.3. Parameter estimation

In this section, we address Problem 5.1.2 and describe the parameter estimation for the model with branching (5.6). This can be easily adapted to the non-branching case. Firstly, we outline how we parameterize the infinitely dimensional rate functions. Then, we develop a general likelihood for the branching case that includes all data types and subsequently explain the details of each term. Lastly, we explain the need for regularization and how the corresponding hyper parameters can be chosen.

5.3.1. Parameterization using splines

The parameters as defined in (5.1) and (5.6) are infinite dimensional. We chose to parameterize these infinite dimensional rates using a finite number of parameters that give the

value of the respective rate at a certain number of equidistant points along the cell state. These values are connected by a natural cubic spline (ncs) (Figure 5.3 C) Luzyanina et al. (2009),

$$\begin{aligned} v_b(s_b) &= \exp(\text{ncs}(s_b|\vec{\alpha}_{v,b})) \\ D_b(s_b) &= \exp(\text{ncs}(s_b|\vec{\alpha}_{D,b})) \\ g_b(s_b) &= \text{ncs}(s_b|\vec{\alpha}_{g,b}). \end{aligned} \tag{5.23}$$

The vectors $\vec{\alpha}_{v,b}$, $\vec{\alpha}_{D,b}$ and $\vec{\alpha}_{g,b}$ are the values of the spline at predefined nodes. The number of nodes might differ between branches and rates. To guarantee positivity of drift and diffusion rates, the spline parameters of v and D are chosen in log-space, the value of the spline at the grid points is evaluated in log-space and then the exponential of these values is computed. The combination of all the unknowns of the model is denoted by $\theta \in \mathbb{R}^{n_\theta}$. For the case of two branches this yields $\theta = (\vec{\alpha}_{D^1}, \vec{\alpha}_{D^2}, \vec{\alpha}_{v^1}, \vec{\alpha}_{v^2}, \vec{\alpha}_{g^1}, \vec{\alpha}_{g^2}, \delta_{12}, \delta_{21})$.

5.3.2. Likelihood

We employed ML estimation, to compute the unknown model parameters from the data. As we used the negative log-likelihood in the optimization, we directly present $-\log(\mathcal{L}(\theta))$ in this section. Assuming independence between the different measurement types, we constructed the negative log-likelihood by regarding each type individually. Firstly, the fit of the simulated cell state number densities to densities obtained from the samples of single-cell cell state measurements, $S_{b,t}^r$ is assessed. We denote the density term by $j_{S,b,t}(\theta)$. This fit is evaluated for each branch $b \in B$ separately and for each single-cell measurement time point $t \in T^{cdf}$. The second part, denoted by $j_{N,t}(\theta)$, describes the fit of the simulated total population size to the measured one and is evaluated for each population size measurement time point $t \in T^N$. The third part of the negative log-likelihood is $j_{w,b,t}(\theta)$. It compares the fit of the branch weights, i.e., the fraction of cells on each branch for each time point $t \in T^{cdf}$. As the branch weights sum up to one, we only considered the fraction of cells on all branches except the last one. The resulting negative log-likelihood consists of a sum of the three terms,

$$-\log(\mathcal{L}(\theta)) = \left(\sum_{b \in B} \sum_{t \in T^{cdf}} j_{S,b,t}(\theta) \right) + \left(\sum_{t \in T^N} j_{N,t}(\theta) \right) + \left(\sum_{b \in B \setminus b_{max}} \sum_{t \in T^{cdf}} j_{w,b,t}(\theta) \right). \tag{5.24}$$

In the following sections, we explain each term in detail.

5.3.2.1. Number density term of the log-likelihood

As the numerical implementation using finite volumes described in Section 5.2.3 simulates the average concentration in a grid interval, $u_i(t)$, rather than the concentration itself, the model can not directly provide the probability to find a cell at a cell state s , $p(s, t)$, but only the probability of finding a cell in an interval $[s_i, s_{i+1}[$, $p_i(t)$, defined in (5.22). Hence, we assessed the fit to the cell state measurements based on the distance between two cumulative density functions (cdf). To improve readability we drop the subscript of s_b in the following equations. Further, we will make the dependence on the parameters θ explicit by including them as an argument. The cdf of the simulated density on branch b at time t is

$$\text{cdf}_{u_b}(\theta, s, t) = \frac{\int_0^s u_b(\theta, \bar{s}, t) d\bar{s}}{\int_0^1 u_b(\theta, \bar{s}, t) d\bar{s}}. \quad (5.25)$$

Using the numerical implementation presented in Section 5.2.3, the cdf is computed by

$$\text{cdf}_{u_b}(\theta, s, t) = \sum_{i=1}^{\lfloor s/h \rfloor + 1} p_{b,i}(\theta, t). \quad (5.26)$$

(It is implicitly assumed, that $p_{2,i}(\theta, t) = 0$ for $i < n_a$.) For the set of single-cell cell state observations, $S_{b,t}^r$, at time point $t \in T^{cdf}$ we computed the empirical cumulative density function (ecdf),

$$\text{ecdf}_{S_{b,t}^r}(s) = \frac{1}{|S_{b,t}^r|} \sum_{s' \in S_{b,t}^r} \mathbb{1}_{s' \leq s}. \quad (5.27)$$

For each branch b and time point t , we also considered the ecdf of the pooled set of cell state samples in all replicates, $S_{b,t} = \bigcup_{r \in R_t} S_{b,t}^r$,

$$\text{ecdf}_{S_{b,t}}(s) = \frac{1}{|S_{b,t}|} \sum_{s' \in S_{b,t}} \mathbb{1}_{s' \leq s}. \quad (5.28)$$

To measure the distance between two distributions, the Kolmogorov-Smirnov (K-S) distance is often used and there also exists a statistical model for the comparison of two samples based on this distance. However, the K-S test derived from the statistic is in many cases too sensitive and not a useful error model for the population distribution. One example where the K-S test was shown to be overly sensitive was in the context of immunofluorescence histogram comparison (Lampariello, 2000). Further, the (two-sided) K-S test tests whether two samples come from the same or different distributions. However, measurement errors, e.g., a shift of the distribution, might change the measured

distribution. It is not easy to determine what and how measurement errors impact the measured cell state distribution in detail and how they could be modeled statistically. We investigate the statistics of the cell state distribution also on experimental data in the application example (Section 5.4.3). As more fitting statistics are not yet available, we chose a non generative error model and assume normally distributed errors on the distance between cdfs. This log-likelihood choice is also capable to incorporate replicates. As distance between two cdfs, we studied the area between the curves,

$$\mathcal{A}(\text{cdf}_1(s), \text{cdf}_2(s)) := \int_0^1 |\text{cdf}_1(s) - \text{cdf}_2(s)| ds, \quad (5.29)$$

i.e., the \mathcal{L}^1 -norm. As this choice leads to a non-smooth optimization problem, one could also use the \mathcal{L}^2 -norm

$$\mathcal{A}_2(\text{cdf}_1(s), \text{cdf}_2(s)) = \left(\int_0^1 (\text{cdf}_1(s) - \text{cdf}_2(s))^2 ds \right)^{\frac{1}{2}}. \quad (5.30)$$

For all time points, we computed the area between the ecdf of the pooled sample, $\text{ecdf}_{S_{b,t}}(s)$, and each of the ecdfs of the individual replicates, $\text{ecdf}_{S_{b,t}^r}(s)$ and then determined the mean,

$$\bar{y}_{b,t}^{\mathcal{A}} = \frac{1}{|R_t|} \sum_{r \in R_t} \mathcal{A} \left(\text{ecdf}_{S_{b,t}^r}(s), \text{ecdf}_{S_{b,t}}(s) \right), \quad (5.31)$$

and the standard deviation,

$$(\sigma_{b,t}^{\mathcal{A}})^2 = \frac{1}{|R_t|} \sum_{r \in R_t} \left(\mathcal{A} \left(\text{ecdf}_{S_{b,t}^r}(s), \text{ecdf}_{S_{b,t}}(s) \right) - \frac{1}{|R_t|} \sum_{r \in R_t} \mathcal{A} \left(\text{ecdf}_{S_{b,t}^r}(s), \text{ecdf}_{S_{b,t}}(s) \right) \right)^2, \quad (5.32)$$

of the area between the curves for the replicates R_t . Further, we calculated the area between the cdf generated by our model and the ecdf of the pooled sample,

$$y_b^{\mathcal{A}}(\theta, t) = \mathcal{A} \left(\text{cdf}_{u_b}(\theta, s, t), \text{ecdf}_{S_{b,t}}(s) \right) = \int_0^1 |\text{cdf}_{u_b}(\theta, \bar{s}, t) - \text{ecdf}_{S_{b,t}}(\bar{s})| d\bar{s}. \quad (5.33)$$

An illustration of the areas is provided in Figure 5.3 D. All in all, we obtained the following negative log-likelihood for the densities,

$$j_{S,b,t}(\theta) = \frac{1}{2} \left(\log(2\pi(\sigma_{b,t}^{\mathcal{A}})^2) + \left(\frac{\bar{y}_{b,t}^{\mathcal{A}} - y_b^{\mathcal{A}}(\theta, t)}{\sigma_{b,t}^{\mathcal{A}}} \right)^2 \right). \quad (5.34)$$

5.3.2.2. Population size term of the log-likelihood

We derived the log-likelihood term for the population size, by assuming normally distributed measurement noise in the population size replicates. However, we did not consider the individual observations, $N_{t,z}$, in the log-likelihood as often only the mean of the replicated population size measurements, \bar{N}_t , is reported. Instead, we considered additive normally distributed noise on \bar{N}_t with mean zero. The standard deviation of this noise is the standard error of the mean of the observations at time point $t \in T^N$,

$$\sigma_t^N = \frac{\sigma_t^{N,obs}}{\sqrt{|Z_t|}} \quad (5.35)$$

where $\sigma_t^{N,obs}$ is the observed standard deviation of the samples at time point t and $|Z_t|$ the number of population size measurement replicates. The part of the negative log-likelihood term corresponding to the population size measurements is given by

$$j_{N,t}(\theta) = \frac{1}{2} \left(\log \left(2\pi (\sigma_t^N)^2 \right) + \left(\frac{\bar{N}_t - N(\theta, t)}{\sigma_t^N} \right)^2 \right), \quad (5.36)$$

where $N(t, \theta)$ is the population size simulated by the model for the parameters θ at time point t across all branches,

$$N(\theta, t) = \sum_{b \in B} \int_0^1 u_b(\theta, s_b, t) ds_b, \quad (5.37)$$

or as calculated in (5.21).

5.3.2.3. Branch weight term of the log-likelihood

As a third part, we added a term to the negative log-likelihood that describes the fit of the fraction of cells on each branch. This addition is necessary, as the information on branch weights is not included in any of the above parts and could otherwise be completely incorrect at the optimal point.

Statistically, if the number of cells is binomially distributed between the two branches, the resulting fraction on a branch is approximately normally distributed. This follows from the central limit theorem (e.g., from the central limit theorem for Bernoulli-sequences (Georgii, 2009, p. 138)). Motivated by this relation, we chose a normally distributed likelihood for the branch weights. In general, for each replicate the fraction of cells on a branch

is a sample from a normal distribution with the true success probability as expectation. Depending on the sample size, the standard deviation varies. However, in the case of single-cell experiments, other sources of measurement error than the sampling error modeled by the binomial distribution can occur. Therefore, we decided to estimate the effective standard deviation from the data. This also ensures consistency with the other parts of the negative log-likelihood. To stay in agreement with the second part of the negative log-likelihood, we consider the mean of the branch weights over the replicates for measurement time points $t \in T^{cdf}$,

$$\bar{w}_{b,t} = \frac{1}{|R_t|} \sum_{r \in R_t} w_{b,t}^r. \quad (5.38)$$

Accordingly, the standard deviation is the standard error of the mean of the observations at time point t ,

$$\sigma_{b,t}^w = \frac{\sigma_{b,t}^{w,obs}}{\sqrt{|R_t|}} \quad (5.39)$$

where $\sigma_{b,t}^{w,obs}$ is the observed standard deviation of the samples at time point t and $|R_t|$ the number of replicates of cell state distribution observations at time point t and branch b . The weight of branch b at time point t simulated by the model for the parameters θ is calculated as,

$$w_b(\theta, t) = \frac{\int_0^1 u_b(\theta, s, t) ds}{\sum_{\bar{b} \in B} \int_0^1 u_{\bar{b}}(\theta, s, t) ds} \quad (5.40)$$

or using the numerical simulation

$$w_b(\theta, t) = \frac{N_b(\theta, t)}{N(\theta, t)}, \quad (5.41)$$

with $N_b(\theta, t)$ and $N(\theta, t)$ as calculated in (5.21). Consequently, the negative log-likelihood of the fraction of cells on a branch b given the model is given by

$$j_{w,b,t}(\theta) = \frac{1}{2} \left(\log(2\pi(\sigma_{b,t}^w)^2) + \left(\frac{\bar{w}_{b,t} - w_b(\theta, t)}{\sigma_{b,t}^w} \right)^2 \right). \quad (5.42)$$

5.3.3. Regularization

The parametrization of the model rates in terms of natural cubic splines is potentially very flexible and can lead to very rapidly varying parameter values if the estimated values at the spline nodes, e.g., $\vec{\alpha}_v$, differ a lot between neighboring nodes. This variability can lead to overfitting and identifiability problems. Hence, we used a quadratic regularization on

the difference between parameters at two neighboring spline nodes, e.g., $(\vec{\alpha}_D)_1 - (\vec{\alpha}_D)_2$, in addition to the likelihood terms (5.24). The combination leads to the following objective function for the parameter estimation,

$$J_\rho(\theta) = -\log(\mathcal{L}(\theta)) + \rho \left(\sum_{b \in B} \left[\sum_{i=1}^{n_b^D-1} ((\vec{\alpha}_{D^b})_{i+1} - (\vec{\alpha}_{D^b})_i)^2 + \sum_{i=1}^{n_b^v-1} ((\vec{\alpha}_{v^b})_{i+1} - (\vec{\alpha}_{v^b})_i)^2 + \sum_{i=1}^{n_b^g-1} ((\vec{\alpha}_{g^b})_{i+1} - (\vec{\alpha}_{g^b})_i)^2 \right] \right), \quad (5.43)$$

where n_b^D , n_b^v and n_b^g are the number of spline nodes on branch $b \in B$ for the splines $D^b(s)$, $v^b(s)$ and $g^b(s)$ on each branch, respectively. The regularization hyper-parameter, ρ , determines the strength of the regularization. How ρ was chosen is explained in the following section.

The regularization can also be interpreted as a prior in a Bayesian sense,

$$-\log(p(\theta)) = \rho \left(\sum_{b \in B} \left[\sum_{i=1}^{n_b^D-1} ((\vec{\alpha}_{D^b})_{i+1} - (\vec{\alpha}_{D^b})_i)^2 + \sum_{i=1}^{n_b^v-1} ((\vec{\alpha}_{v^b})_{i+1} - (\vec{\alpha}_{v^b})_i)^2 + \sum_{i=1}^{n_b^g-1} ((\vec{\alpha}_{g^b})_{i+1} - (\vec{\alpha}_{g^b})_i)^2 \right] \right). \quad (5.44)$$

The normalization is included in ρ . In this interpretation, the negative log-likelihood with regularization is a negative log-posterior, $J_\rho(\theta) = -\log(p(\theta|D))$, with $p(\theta|D) = \mathcal{L}(\theta)p(\theta)$.

5.3.4. Selection of the regularization hyper-parameter

We selected the regularization hyper-parameter using leave-one-out cross-validation on a time point basis. We removed the complete set of data points corresponding to one time point and fitted the model on the reduced data set. Then, we evaluated the negative log-likelihood without regularization on the data at the withheld time point. This leave-one-out fitting was performed for each time point (except the initial time point $t = 0$). Adding up the negative log-likelihood values without regularization for each withheld time point yielded a score based on which the hyper-parameter was selected as the one minimizing the score.

5.4. Application to T-cell maturation data

In this section, we use the continuous cell state model with branching and the parameter estimation approach presented above for the analysis of a data set of T-cell maturation. Before the application to experimental data, we evaluated our approach on various cases of simulated data. These simulation studies can be found in the manuscript (Fischer et al., 2019).

T-cells are an important part of the adaptive immune system and play a relevant role in many functions and diseases. During embryogenesis T-cells develop from hematopoietic progenitors and mature in the thymus. The developmental stages are mainly characterized by expression of CD4 and CD8. During differentiation cells have to pass check points and selection to ensure correct biological function. In the following, we describe the experiment, data processing and the specific details of the numerical implementation and analyze the parameter estimation results.

5.4.1. Experiment and data processing

The experiments were performed by the biological partner, the group of Rene Maehr at the University of Massachusetts. The analysis with diffusion maps, the cell state extraction and the detection of the branching region together with analyses on the expression data was mainly performed by David Fischer and is presented in detail in (Fischer et al., 2019).

In the experiment, single-cell sequencing samples were collected from mouse thymus using a Drop-seq protocol. The samples were collected daily from 12.5 to 19.5 days after fertilization: 12.5 (3), 13.5 (3), 14.5 (2), 15.5 (2), 16.5 (3), 17.5 (2), 18.5 (2), 19.5 (2). The number of replicates is indicated in brackets. The population size measurements were obtained for the same time points, but from independent individuals. The measurements for the days 12.5-17.5 after fertilization using 5 replicates each were taken from published data (Cook, 2010) and augmented with two replicates of population size measurements performed by the group of Maehr at days 18.5 and 19.5 after fertilization. In the mathematical model, we shifted all values by 12.5 and consider the model time frame $t = 0$ to $t = 7$ in the following.

The cell state was extracted using a diffusion map on the combination of all single-cell sequencing data and diffusion pseudo time. The scatter plot of the first two diffusion components is shown in Figure 5.4 A. Already looking at the first two diffusion components branching is clearly evident. By reference to the expression profiles two distinct cell

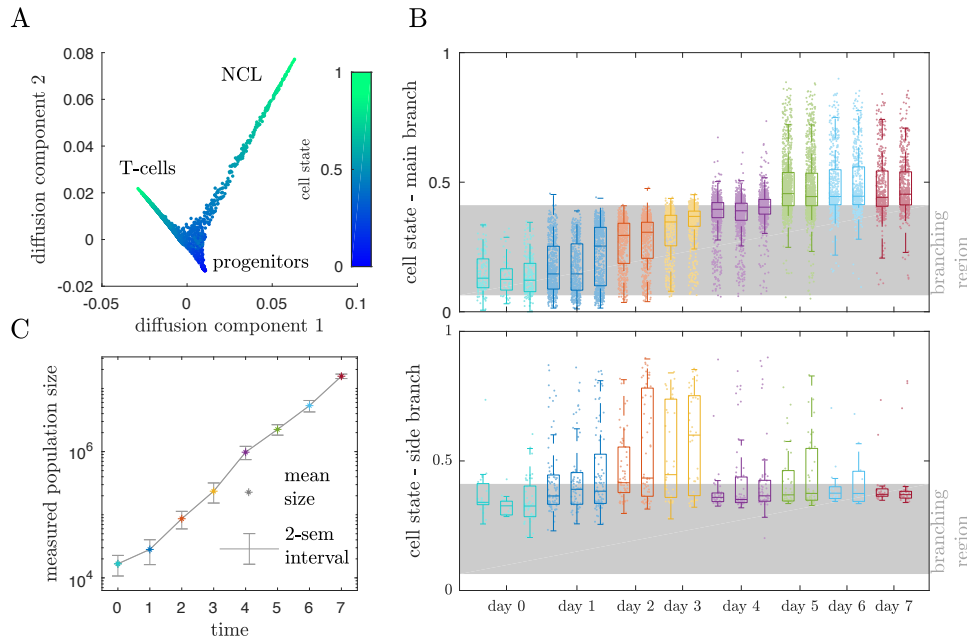


Figure 5.4.: T-cell maturation data. (A) Scatter plot of the first two diffusion components of the diffusion map for the single-cell T-cell maturation data. The color denotes the cell state assigned to the cell using diffusion pseudo time. (B) Cell state distribution for each replicate and each day on the main and side branch, where the color indicates the respective day. The branching region is marked in grey. (C) Mean population size and 2-times standard error of mean interval of the population size measurements.

lineages were identified in the data, T-cells and non-conventional lymphocytes (NCLs) which is in concordance with previous findings. Molecularly, these two lineages correspond to the two branches in the diffusion map. We assigned the progenitor up to T-cell lineage to the main branch and the NCL lineage to the side branch. Figure 5.4 B depicts the cell state distribution on the main and the side branch for each day and replicates as dots and box plots. The mean population size is displayed in Figure 5.4 C. Around the mean population size we plotted a 2-standard error of mean interval. It stands out that during the last three time points the distribution does not change (Figure 5.4 B) and seems to reach a steady state. However, the overall population size is still increasing.

5.4.2. Numerical implementation

As the data processing revealed two branches in the data, we used model (5.6). For the numerical simulation of the PDE, we chose a discretization into 299 intervals on the main branch. The branching region extended from interval $n_a = 20$ to interval $n_e = 123$. The kernel density estimation in the computation of the initial condition (5.10) was performed

Table 5.1.: Lower bounds, upper bounds and scale used for fitting of the parameters of model for T-cell maturation.

Parameter name	Lower bound	Upper bound	Scale used for optimization
D_i	$10^{-4.5}$	10^0	log
v_i	10^{-5}	10^0	log
g_i	-6	5	lin
δ_{ij}	$10^{-4.5}$	10^2	log

using the Matlab function `ksdensity.m`. The resulting ODE was solved numerically using the Sundials CVODE suite and AMICI (Fröhlich et al., 2017b) as Matlab interface. We chose nine nodes on the main branch and three nodes on the side branch for the natural cubic spline parametrization of the rates ($D_b(s)$, $v_b(s)$ and $g_b(s)$) (Section 5.3.1), yielding the parameter vector

$$\theta = ((\vec{\alpha}_{D^1})_1, \dots, (\vec{\alpha}_{D^1})_9, (\vec{\alpha}_{D^2})_1, \dots, (\vec{\alpha}_{D^2})_3, (\vec{\alpha}_{v^1})_1, \dots, (\vec{\alpha}_{v^1})_9, (\vec{\alpha}_{v^2})_1, \dots, (\vec{\alpha}_{v^2})_3, (\vec{\alpha}_{g^1})_1, \dots, (\vec{\alpha}_{g^1})_9, (\vec{\alpha}_{g^2})_1, \dots, (\vec{\alpha}_{g^2})_3, \delta_{12}, \delta_{21}). \quad (5.45)$$

To ensure no drift at the right-hand boundary, i.e., $v_b(1, t) = 0$, without causing numerical instabilities in the simulation as well as in the spline interpolation of v_b , we linearly rescaled the observations to the interval $[0, 0.9]$ and ran simulations on the interval $[0, 1]$. On the interval $[0.9, 1]$ (more precisely the control intervals $i = 270$ to $i = 299$) the drift was reduced to 0 by a smooth hermite c-spline. Values in the interval $[0.9, 1]$ correspond to developmental states which were not observed and simulations with considerable cell mass in this interval are unfavorable in the optimization.

For the optimization, we used a multi-start approach with gradient information for the optimization as implemented in the PESTO-Toolbox (Stapor et al., 2018b). Initial parameters for multi-start local optimization were sampled using latin-hypercube sampling. The parameter bounds for sampling and optimization are presented in Table 5.1. We performed up to 120 multi-starts using an interior point algorithm with BFGS. If a plateau of 6 starts was achieved at 40 or 80 multi-starts we did not perform the full 120 starts. The confidence intervals were computed using profile posteriors (Raue et al., 2013a) as implemented in the PESTO-Toolbox.

5.4.3. The K-S distance is overly sensitive for the T-cell data

In Section 5.3.2.1, we touched on how the K-S test is often overly sensitive. Hence, we evaluated this claim for the T-cell differentiation data.

Firstly, we tested all replicates of the T-cell development data from the same time point against each other using a two sample K-S test. Of the 14 possible combinations seven were judged to be not coming from the same distribution at a significance level of $\alpha = 0.05$. Hence, based on a K-S test the samples for several time points would not be assumed to come from the same distribution (Figure 5.5 A).

To further study the variability in the data, we generated a distribution of K-S distances by sampling from the same underlying distribution. We drew 1000 sub-samples from the distribution of the pooled data corresponding to a day t , $S_{b,t}$, and evaluated the K-S distance and the area between the sampled and the pooled cdf. As the number of measured single cells in the T-cell data varied between replicates, we tried to recreate this in the generated data by sampling the sizes of the 1000 sub-samples from a distribution with the same first two moments as the measured samples. For days 1 to 8, we sampled from a normal distribution with the same mean and variance as the single-cell sample sizes. For day 0, the sub-sample size was sampled from a log-normal distribution with the same mean and variance as the single-cell sample sizes. When comparing the distance of the measured and the pooled sample to the created distribution of distances, we found that for several days the measurements lie outside the distribution indicating that the measurements are unlikely just samples from the pooled distribution (Figure 5.5 B). The K-S distance did not capture the full variability between replicates.

5.4.4. Model fit and regularization

We fit the continuous model (5.6) to the T-cell maturation data and found that it can describe the data (Figure 5.6 A and B). The statistical model we developed uses regularization to avoid overfitting. Therefore, the fit for regularization hyper-parameter $\rho = 10$ depicted in the figure shows some trade off between fit and parameter consistency. To select the regularization hyper-parameter ρ , we investigated the regularizations $\rho = \{0, 1, 10, 30, 100, 300, 1000\}$ in a leave-one-time-point-out cross validation (Section 5.3.4). Using the parameter estimates obtained on the reduced data sets, the log-likelihood on the data from the left-out time point can be computed. The leave-one-time-point-out cross validation suggested an intermediate to small regularization strength

5.4. Application to T-cell maturation data

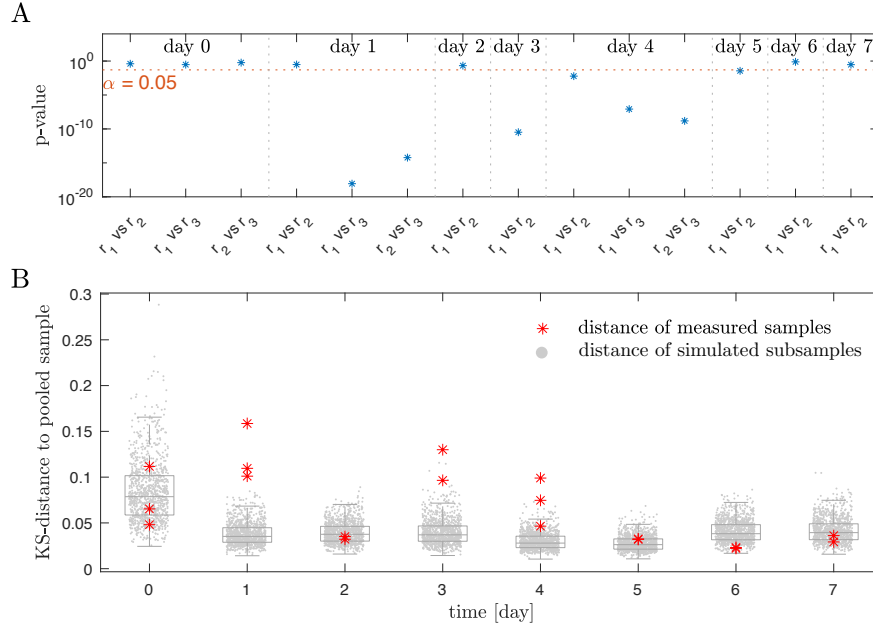


Figure 5.5.: Evaluation of the K-S distance as statistical model for T-cell maturation data. (A) p-values for pairwise K-S tests of single-cell cell state measurement replicates taken at the same time point (blue stars) and confidence level $\alpha = 0.05$ (orange dotted line). If the p-value of a comparison is smaller than $\alpha = 0.05$, there is little evidence that the compared samples were generated from the same underlying distribution. (B) For each measurement day the distribution of K-S distances between the ecdf of the thousand generated subsamples and the ecdf of the pooled sample is depicted as dots and box plots. The distances between the ecdf of the measured and the pooled samples are plotted in red.

of $\rho = 10$, as the prediction error on the left out time point is smallest (Figure 5.6 C). Comparing the contribution of different days to the over all prediction error, shows that some days can consistently be predicted better than others. As an example, we visualized the prediction for day 2 and day 3. The difference in prediction quality also implies that temporal sampling frequency is relevant for imputation of missing time points. While the cross validation favors $\rho = 10$, the best estimates using other regularizations show a similar behavior (Figure 5.6 D). Especially, low drift rates from around $s = 0.3$ to $s = 0.5$ with subsequent increase and regions of positive and negative proliferation rate seem to pertain through ranges of regularization. The uncertainty analysis shows that regions of higher diffusion are more uncertain than regions of small diffusion. Further, the confidence intervals around negative and positive proliferation are still clearly negative respective positive. This consistency and certainty reinforces the validity of biological interpretations.

5.4.5. Parameter estimates identify checkpoints and selection in T-cell maturation

Biologically, we identified different behaviors in the dynamics on the main branch that can be interpreted in terms of T-cell maturation. The advection parameter fit (Figure 5.6 D) uncovers two intervals of high drift around $s = 0.2$ and $s = 0.6$. These regions correspond to rapid transcriptomic development and deterministic behavior of individual cells. On a molecular level, this rapid development could be validated by looking at the regulation and activity of transcription factors. In the cell state interval of $s = 0.3$ to $s = 0.5$, we find a region of low drift and non-zero diffusion leading to non-deterministic development.

Molecular analysis suggests that the developmental checkpoint of β -selection lies in this area. At $s = 0.4$ to $s = 0.7$, after the β -selection area, we find positive proliferation that monotonically decreases with cell state and eventually, at around $s = 0.7$, becomes negative. Biologically, T-cells that pass beta-selection are known to divide rapidly and then undergo positive and negative selection (Yui and Rothenberg, 2014) which coincides with the estimated rates.

5.5. Discussion

In this chapter, we studied estimation of population dynamics of heterogeneous cell populations from high dimensional single-cell data. To reduce the dimensionality of the mRNA sequencing data a one dimensional cell state was extracted and assigned to the individual cells. To address Problem 5.1.1, we developed PDE models describing the differentiation on a population density level as well as the proliferation dynamics of the total population for non-branching and branching processes. We further outlined the numerical simulation. After the publication of this work on bioRxiv, Cho et al. (2018) published a diffusion-advection-reaction model similar to the model presented here also using diffusion pseudo time as cell state. Their work did, however, not include an approach to parameter estimation. In their model, Cho et al. (2018) separated the diffusion term into fluctuations in the cell state due to randomness in transcription and due to the projection on one dimension. In contrast, we combined both contributions into a single diffusion term. In future work, this separation of contributions could be included into the model to investigate if the individual effects can be identified during parameter estimation and if this level of precision is necessary. Further, in this work we implemented a branching process with two terminal cell fates. While not considered here, cases with more than two terminal fates are




Figure 5.6.: Estimation and cross validation results. (A) Fit of the density simulated with the best parameter estimate for regularization $\rho = 10$ (color) and superimposed histograms of the measured cell states for all replicates (gray) resolved by time points. For day 2 and day 3 the predicted density at this time point for the cross validation fold without this time point and regularization $\rho = 10$ is depicted (broken line in dark orange and dark yellow). (B) Population size measurement (gray) and the fit for $\rho = 10$ (black). As an example, also the fits for the best parameter estimate in the cross validation folds without day 2 and day 3 (broken line in dark orange and dark yellow) are depicted. The predicted value at day 2 and day 3 is marked with a star on the corresponding curve. (C) The prediction error score of the regularization strengths $\rho \in \{1, 10, 30, 100, 300, 1000\}$ is visualized as the height of the bar. The contribution of individual leave-one-time-point-out cross validation folds is characterized by the height of the colored segments (color coding like (A)). Due to their minimal contribution, some folds are not or barely visible, e.g., day 1. (D) Estimated rates on the main branch obtained by optimization with regularization for $\rho = 10$ and the 99% confidence intervals at the parameters $((\vec{\alpha}_{D^1})_1, \text{etc.})$. The profile likelihood computation of the confidence interval of $(\vec{\alpha}_{v^1})_1$ yielded no upper bound inside the parameter bounds and is set to one for the illustration.

relevant in many developmental processes especially when studied for a longer time frame, e.g., hematopoiesis. These cases could be modeled by coupling more branches. However, we expect the discretization using finite volumes to become prohibitive, as the size of the resulting discretized PDE increases. Developing an alternative model description of pro-

cesses with more than two branches on the one hand and accelerating the simulation of the PDE model on the other hand are open points for further research.

To address Problem 5.1.2, we developed a parameter estimation framework including a parameterization of the infinitely dimensional parameters and a likelihood. We then applied the model in the study of T-cell maturation data. We found that the model could describe the data and that leave-one-time-point-out cross validation suggested an intermediate regularization strength. The estimated rates and their interpretation uncover known biological aspects without directly incorporating knowledge about the driving molecular processes in the estimation. We chose to base the likelihood on the assumption of normally distributed measurement noise. This error model is not generated by a statistical model of the data generation and measurement procedure as this is not available. Hence, the underlying statistics need to be studied more in further experiments to develop such a generative model and the likelihood can then be adapted in the light of these new findings. We further, used an \mathcal{L}^1 distance between densities in the formulation of the likelihood describing the single-cell cell state measurements. This yields a non-smooth optimization problem that is currently addressed using interior point optimization with BFGS. The optimization in the current implementation is comparably slow and expensive. Choosing the \mathcal{L}^2 distance and considering the corresponding smooth problem might accelerate the parameter estimation. It is an interesting open question how the convergence behavior of the smooth problem differs. A general comparison between these alternative likelihoods is a point of future research.

Chapter 6.

Conclusion and outlook

Increasing model size and data dimensions pose big challenges for the estimation of model parameters for biological models. With more detailed systems biology and systems medicine applications on the one hand and with advancing experimental techniques on the other it is likely that model and data dimensions will continue to increase. Hence, the need to develop efficient and robust estimation techniques arises. One way to approach this is to reduce the dimension of the problem. In this thesis, two different strategies to reduce the dimensionality of parameter estimation problems by exploiting structures in both, the parameter space and in the data, were investigated.

For the first dimension reduction approach developed in Chapters 3 and 4, we considered ODE and PDE constrained estimation with additional steady-state constraints. Steady states are common in biological applications due to, among others, fast calibration of processes or homeostasis. Including the steady-state information into the parameter estimation leads to additional equality constraints for the optimization. If no analytical solution for the steady state as function of the parameters is available, the optimization variables include the parameter and the corresponding steady states. Hence, constraining the search space to the steady-state manifold, i.e., the admissible parameter-state combinations, reduces the dimension of the problem. The methods we developed in Chapter 3 for ODE models and in Chapter 4 for PDE models exploit the first order structure of the steady-state manifold to choose a state update on the steady-state manifold for a given parameter update. For ODEs we constructed two methods that incorporated ideas from retraction mappings and continuous analogues, respectively. In Chapter 4, we extended the continuous analogue approach to PDE models. For the continuous analogue methods, we proved stability of optimal parameter-state pairs for both ODE and PDE constrained optimization using singular perturbation theory in the ODE case and in Lyapunov functions in the PDE case. Additionally, singular perturbation theory provided insight into

Chapter 6. Conclusion and outlook

the general behavior of the continuous analogue, i.e., first moving to the steady-state manifold and then following it to the optimal parameter. In several applications to simulated and real data, we illustrated that the new methods outperformed standard solvers for constrained optimization in terms of robustness and time efficiency. The applications also demonstrated that if the analytical steady state is available solving the unconstrained optimization problem is preferable. Both proofs for asymptotic stability show that the value of the retraction factor λ needs to be chosen sufficiently large. Hence, we illustrated the influence of the retractions factor on runtime and convergence and investigated possible heuristics for the retraction factor choice.

Ideas for future work (see also Sections 3.4 and 4.5) include the extension of the proof to settings with non-identifiable parameters – based on the application results, we assume that the continuous analogue should usually return a point in the non-identifiable subspace – as well as the development and implementation of an automated and adaptive retraction factor choice. For optimization problems with parabolic and elliptic PDE constraints, we addressed the possibility of including the solution of the parabolic PDE as additional state variable thus eliminating the need for a solution operator for the PDE. Further, we discussed the application to problems with infinite dimensional parameters. Beyond these extensions, the acceleration of optimization achieved by the novel methods can also be employed to increase the efficiency of practical identifiability analysis and uncertainty quantification by speeding up the optimization steps in the profile likelihood calculation (Raue et al., 2009). In recent studies, a dynamical systems approach to profile likelihood calculation based on ideas by Chen and Jennrich (2002) was studied and shown to perform favorably (Boiger et al., 2016; Stapor et al., 2018a). A combination of the continuous analogue approach to steady-state constrained optimization with the dynamical system for profile likelihood computation could potentially further increase the efficiency of likelihood profile calculations for steady-state constrained problems. While we did not explicitly address this in Chapter 4, considering a multi experiment setting for PDEs similarly to the one Chapter 3 should be straight forward to implement. Lastly, we only considered the log-likelihood for normal and log-normally distributed measurement noise as objective function in our application. However, the proof did not assume a specific objective function and an application of the methods to other objective function types might be interesting.

The second approach for dimension reduction focused on reducing the dimension of the data by extracting the most relevant dimensions using structure inherent in the high-dimensional measurements (Chapter 5). In our case high-dimensional gene expression data was reduced to a one-dimensional cell state space. On this reduced space we developed a

dynamic PDE model that describes population dynamics of a heterogeneous population. The model includes differentiation and proliferation dynamics and yields an advection-diffusion-reaction equation for the evolution of the number density over the cell state. We developed a non-branching model for differentiation into a single cell-type and a branching model that describes the differentiation into two cell-types. For the simulation of the PDE, we implemented a finite volume discretization together with the method of lines to accurately account for the population growth dynamics. The cell state dependent infinite dimensional rates of the PDE were parameterized using splines to yield a set of finite dimensional model parameters for estimation. We further developed a statistical likelihood framework for parameter estimation that is capable of incorporating repeated single-cell RNA sequencing experiments and population size measurements. We applied the model and parameter estimation framework to single-cell RNA sequencing data for T-cell differentiation in thymus and were able to identify known biological mechanisms without having incorporating these as prior knowledge.

In Section 5.5, we outlined some possible extensions to our approach: The branching model is currently only implemented for two branches, however, in principle more cell fates can be incorporated. Including more branches will likely scale up the computation time of the current implementation. Hence, an efficient implementation would be required. We further outlined the possibility of considering other likelihoods and possibly developing a generative statistical model. In a further step the uncertainty in the dimension reduction to the cell state could be incorporated into the likelihood. Beyond this, future work might include an extension of the model to more than one cell state dimension. With this the branching to different cell states could be modeled in a two dimensional space and the cell fate decision would be included in the advection term. In addition to single-cell RNA sequencing data and population size, more data types, e.g., single cell Assay for Transposase-Accessible Chromatin using sequencing (ATAC-seq) or protein expression measurements, could be included in the parameter estimation. While RNA sequencing determines the cell state at the level of mRNA expression, ATAC-seq determines it at the chromatin level and protein expression at a functional level. Hence, the additional inclusion of these data types could provide a deeper understanding of the cell state trajectory. Finally, including cell age into structured population models has been shown to describe proliferation dynamics more accurately as it allows for a delay in cell division times (Hross and Hasenauer, 2016). Hence, future work might consist of incorporating an age or cell cycle variable into the PDE and describing cell division and cell death in dependence on cell age or cell cycle state.

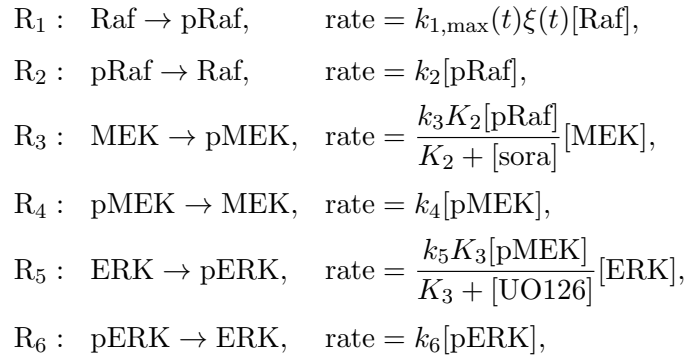
Chapter 6. Conclusion and outlook

In conclusion, this thesis proposed two dimension reduction approaches to facilitate modeling and parameter estimation in high dimensional parameter and data spaces. Furthermore, we illustrated the properties of the developed methods and their applicability using examples of simulated and biological data. Employing these tailored methods can then enable the analysis of biological processes using extensive models and large data sizes and facilitate a move towards a more comprehensive understanding of the underlying biological mechanisms.

Appendix A.

Analytical steady state for the Raf/MEK/ERK signaling model

In Section 3.3.4, we considered a model for Raf/MEK/ERK signaling after release from S-phase arrest considering the following reactions,



with

$$k_{1,\max}(t) = k_{1,0} + k_{1,1} \left(1 - e^{-\frac{t}{\tau_1}}\right) e^{-\frac{t}{\tau_2}}.$$

The model includes inhibition by Sorafenib and UO126 (R_4 and R_6) modeled by a reduction in the kinase activity of pRaf and pMEK. We considered two model hypotheses for the effect of pERK on the phosphorylation of Raf:

H1 Inhibition of Raf phosphorylation by pERK: $\xi(t) = \frac{K_1}{K_1 + [\text{pERK}]}$

H2 No inhibition: $\xi(t) = 1$

After the computations and parametrization outlined in Section 3.3.4.2, we arrived at the following model,

$$\frac{dx_1}{dt} = k_{1,\max}(t)\xi(t)(1 - x_1) - k_2 x_1$$

Appendix A. Analytical steady state for the Raf/MEK/ERK signaling model

$$\begin{aligned}\frac{dx_2}{dt} &= \frac{k_3[\text{Raf}]_0 K_2 x_1}{K_2 + [\text{sora}]} (1 - x_2) - k_4 x_2 \\ \frac{dx_3}{dt} &= \frac{k_5[\text{MEK}]_0 K_3 x_2}{K_3 + [\text{UO126}]} (1 - x_3) - k_6 x_3 \\ y_{1,b} &= s_{1,b}[\text{MEK}]_0 x_2 \\ y_{2,b} &= s_{2,b}[\text{ERK}]_0 x_3\end{aligned}$$

with blot index $b = 1, \dots, 4$. With the definitions $\tilde{K}_1 = K_1/[\text{ERK}]_0$, $\tilde{k}_3 = k_3[\text{Raf}]_0$ and $\tilde{k}_5 = k_5[\text{MEK}]_0$, the steady state for H1 is given by

$$\begin{aligned}\varphi_{0,1}(\theta) &= \left(\tilde{K}_1 k_{1,0} + \left(\tilde{K}_1^2 k_{1,0}^2 + \frac{2\tilde{K}_1^2 k_6 k_{1,0}^2}{\tilde{k}_5} + \frac{\tilde{K}_1^2 k_6^2 k_{1,0}^2}{\tilde{k}_5^2} + \frac{\tilde{K}_1^2 k_4^2 k_6^2 (k_{1,0} + k_2)^2}{(\tilde{k}_3 \tilde{k}_5)^2} + \right. \right. \\ &\quad \left. \frac{2\tilde{K}_1^2 k_4 k_6^2 k_{1,0} (k_{1,0} + k_2)}{\tilde{k}_3 \tilde{k}_5^2} + \frac{2\tilde{K}_1^2 k_4 k_6 k_{1,0} (k_{1,0} + k_2)}{\tilde{k}_3 \tilde{k}_5} + \frac{4\tilde{K}_1 k_2 k_4 k_6 k_{1,0}}{\tilde{k}_3 \tilde{k}_5} \right)^{\frac{1}{2}} + \\ &\quad \left. \frac{\tilde{K}_1 k_6 k_{1,0}}{\tilde{k}_5} - \frac{\tilde{K}_1 k_4 k_6 (k_{1,0} + k_2)}{\tilde{k}_3 \tilde{k}_5} \right) / \\ &\quad \left(2 \left(k_2 + \tilde{K}_1 k_{1,0} + \tilde{K}_1 k_2 + \frac{\tilde{K}_1 k_2 k_6}{\tilde{k}_5} + \frac{\tilde{K}_1 k_6 k_{1,0}}{\tilde{k}_5} \right) \right) \\ \varphi_{0,2}(\theta) &= \left(\left(\tilde{K}_1^2 k_{1,0}^2 + \frac{2\tilde{K}_1^2 k_6 k_{1,0}^2}{\tilde{k}_5} + \frac{\tilde{K}_1^2 k_6^2 k_{1,0}^2}{\tilde{k}_5^2} + \frac{\tilde{K}_1^2 k_4^2 k_6^2 (k_{1,0} + k_2)^2}{(\tilde{k}_3 \tilde{k}_5)^2} + \right. \right. \\ &\quad \left. \frac{2\tilde{K}_1^2 k_4 k_6^2 k_{1,0} (k_{1,0} + k_2)}{\tilde{k}_3 \tilde{k}_5^2} + \frac{2\tilde{K}_1^2 k_4 k_6 k_{1,0} (k_{1,0} + k_2)}{\tilde{k}_3 \tilde{k}_5} + \frac{4\tilde{K}_1 k_2 k_4 k_6 k_{1,0}}{\tilde{k}_3 \tilde{k}_5} \right)^{\frac{1}{2}} + \\ &\quad \left. \tilde{K}_1 k_{1,0} + \frac{\tilde{K}_1 k_6 k_{1,0}}{\tilde{k}_5} - \frac{\tilde{K}_1 k_2 k_4 k_6}{\tilde{k}_3 \tilde{k}_5} - \frac{\tilde{K}_1 k_4 k_6 k_{1,0}}{\tilde{k}_3 \tilde{k}_5} \right) / \\ &\quad \left(\left(\tilde{K}_1^2 k_{1,0}^2 + \frac{2\tilde{K}_1^2 k_6 k_{1,0}^2}{\tilde{k}_5} + \frac{\tilde{K}_1^2 k_6^2 k_{1,0}^2}{\tilde{k}_5^2} + \frac{\tilde{K}_1^2 k_4^2 k_6^2 (k_{1,0} + k_2)^2}{(\tilde{k}_3 \tilde{k}_5)^2} + \right. \right. \\ &\quad \left. \frac{2\tilde{K}_1^2 k_4 k_6^2 k_{1,0} (k_{1,0} + k_2)}{\tilde{k}_3 \tilde{k}_5^2} + \frac{2\tilde{K}_1^2 k_4 k_6 k_{1,0} (k_{1,0} + k_2)}{\tilde{k}_3 \tilde{k}_5} + \frac{4\tilde{K}_1 k_2 k_4 k_6 k_{1,0}}{\tilde{k}_3 \tilde{k}_5} \right)^{\frac{1}{2}} + \\ &\quad \left. \tilde{K}_1 k_{1,0} + \frac{\tilde{K}_1 k_6 k_{1,0}}{\tilde{k}_5} + \frac{k_2 k_4}{1 \tilde{k}_3} \left(2\tilde{K}_1 + \frac{\tilde{K}_1 k_6}{\tilde{k}_5} + 2 \right) + \frac{\tilde{K}_1 k_4 k_{1,0}}{\tilde{k}_3} \left(\frac{k_6}{\tilde{k}_5} + 2 \right) \right)\end{aligned}$$

$$\begin{aligned}
\varphi_{0,3}(\theta) = & \left(\left(\tilde{K}_1^2 (k_{1,0})^2 + \frac{2\tilde{K}_1^2 k_6 k_{1,0}^2}{\tilde{k}_5} + \frac{\tilde{K}_1^2 k_6^2 k_{1,0}^2}{\tilde{k}_5^2} + \frac{\tilde{K}_1^2 k_4^2 k_6^2 (k_{1,0} + k_2)^2}{(\tilde{k}_3 \tilde{k}_5)^2} + \right. \right. \\
& \left. \frac{2\tilde{K}_1^2 k_4 k_6^2 k_{1,0} (k_{1,0} + k_2)}{\tilde{k}_3 \tilde{k}_5^2} + \frac{2\tilde{K}_1^2 k_4 k_6 k_{1,0} (k_{1,0} + k_2)}{\tilde{k}_3 \tilde{k}_5} + \frac{4\tilde{K}_1 k_2 k_4 k_6 k_{1,0}}{\tilde{k}_3 \tilde{k}_5} \right)^{\frac{1}{2}} \\
& + \tilde{K}_1 k_{1,0} + \frac{\tilde{K}_1 k_6 k_{1,0}}{\tilde{k}_5} - \frac{\tilde{K}_1 k_2 k_4 k_6}{1\tilde{k}_3 \tilde{k}_5} - \frac{\tilde{K}_1 k_4 k_6 k_{1,0}}{\tilde{k}_3 \tilde{k}_5} \Big) / \left(\left(\frac{k_6}{\tilde{k}_5} + 1 \right) \right. \\
& \left. \left(\tilde{K}_1^2 k_{1,0}^2 + \frac{2\tilde{K}_1^2 k_6 k_{1,0}^2}{\tilde{k}_5} + \frac{\tilde{K}_1^2 k_6^2 k_{1,0}^2}{\tilde{k}_5^2} + \frac{\tilde{K}_1^2 k_4^2 k_6^2 (k_{1,0} + k_2)^2}{(\tilde{k}_3 \tilde{k}_5)^2} + \frac{4\tilde{K}_1 k_2 k_4 k_6 k_{1,0}}{\tilde{k}_3 \tilde{k}_5} + \right. \right. \\
& \left. \frac{2\tilde{K}_1^2 k_4 k_6^2 k_{1,0} (k_{1,0} + k_2)}{\tilde{k}_3 \tilde{k}_5^2} + \frac{2\tilde{K}_1^2 k_4 k_6 k_{1,0} (k_{1,0} + k_2)}{\tilde{k}_3 \tilde{k}_5} \right)^{\frac{1}{2}} + \tilde{K}_1 k_{1,0} \left(\frac{k_6}{\tilde{k}_5} + 1 \right)^2 \\
& \left. + \frac{k_2 k_4 k_6}{\tilde{k}_3 \tilde{k}_5} \left(\tilde{K}_1 + \frac{\tilde{K}_1 k_6}{\tilde{k}_5} + 2 \right) + \frac{\tilde{K}_1 k_6 k_6 k_{1,0}}{\tilde{k}_3 \tilde{k}_5} \left(\frac{k_6}{\tilde{k}_5} + 1 \right) \right).
\end{aligned}$$

The steady state for H2 is given by

$$\begin{aligned}
\varphi_{0,1}(\theta) &= \frac{k_{1,0}}{k_{1,0} + k_2}, & \varphi_{0,2}(\theta) &= \frac{\tilde{k}_3 \frac{k_{1,0}}{k_{1,0} + k_2}}{\tilde{k}_3 \frac{k_{1,0}}{k_{1,0} + k_2} + k_4} \\
\varphi_{0,3}(\theta) &= \left(\tilde{k}_5 \frac{\tilde{k}_3 \frac{k_{1,0}}{k_{1,0} + k_2}}{\tilde{k}_3 \frac{k_{1,0}}{k_{1,0} + k_2} + k_4} \right) / \left(\tilde{k}_5 \frac{\tilde{k}_3 \frac{k_{1,0}}{k_{1,0} + k_2}}{\tilde{k}_3 \frac{k_{1,0}}{k_{1,0} + k_2} + k_4} + k_6 \right).
\end{aligned}$$

Appendix B.

Finite volume discretization for the branching model

In Section 5.2.3, we considered the finite volume discretization of the branching model on the grid ih for $i = 0, \dots, n_V$ and interval length h . On the main branch the grid points are denoted as $s_{i-\frac{1}{2}} = (i-1)h$ and $s_{i+\frac{1}{2}} = ih$. The center of the interval $[s_{i-\frac{1}{2}}, s_{i+\frac{1}{2}}]$ is denoted with $s_i = ih - \frac{1}{2}h$. If we denote the interval that contains the start of the branching region, s_a , with n_a , i.e., $s_a \in [s_{n_a-\frac{1}{2}}, s_{n_a+\frac{1}{2}}]$, the side branch contains the intervals indexed by n_a, \dots, n_V and the number of intervals on the side branch is $n_V - n_a + 1$. The interval at the end of the branching region is denoted with n_e . The average cell density in an interval $[s_{i-\frac{1}{2}}, s_{i+\frac{1}{2}}]$ at time point t is denoted by $u_{1,i}(t)$ for $u_1(s, t)$ and $u_{2,i}(t)$ for $u_2(s, t)$. We denote the diffusion rate evaluated at $s_{i\pm\frac{1}{2}}$ with $D_{b,i\pm\frac{1}{2}} = D_b(s_{i\pm\frac{1}{2}}, t)$, the advection rate evaluated at $s_{i\pm\frac{1}{2}}$ with $v_{b,i\pm\frac{1}{2}} = v_b(s_{i\pm\frac{1}{2}}, t)$ and the proliferation rate evaluated at s_i with $g_{b,i} = g_b(s_i, t)$ where $b \in \{1, 2\}$ indicates the corresponding branch. The complete ODE resulting from the finite volume discretization of the branching model (5.6) is given by

$$\begin{aligned} \frac{du_{1,1}}{dt} &= -\frac{1}{h^2} D_{1,1+\frac{1}{2}} (u_{1,1} - u_{1,2}) - \frac{1}{2h} v_{1,1+\frac{1}{2}} (u_{1,1} + u_{1,2}) + g_{1,1} u_{1,1} \\ \frac{du_{1,i}}{dt} &= \frac{1}{h^2} \left(D_{1,i-\frac{1}{2}} (u_{1,i-1} - u_{1,i}) - D_{1,i+\frac{1}{2}} (u_{1,i} - u_{1,i+1}) \right) \\ &\quad + \frac{1}{2h} \left(v_{1,i-\frac{1}{2}} (u_{1,i-1} + u_{1,i}) - v_{1,i+\frac{1}{2}} (u_{1,i} + u_{1,i+1}) \right) + g_{1,i} u_{1,i} \\ &\text{for } i \in \{2, \dots, n_V - 1\} \setminus \{n_a, \dots, n_e\} \\ \frac{du_{1,i}}{dt} &= \frac{1}{h^2} \left(D_{1,i-\frac{1}{2}} (u_{1,i-1} - u_{1,i}) - D_{1,i+\frac{1}{2}} (u_{1,i} - u_{1,i+1}) \right) \\ &\quad + \frac{1}{2h} \left(v_{1,i-\frac{1}{2}} (u_{1,i-1} + u_{1,i}) - v_{1,i+\frac{1}{2}} (u_{1,i} + u_{1,i+1}) \right) \\ &\quad + g_{1,i} u_{1,i} - \delta_{12} u_{1,i} + \delta_{21} u_{2,n_a} \\ &\text{for } i \in \{n_a, \dots, n_e\} \end{aligned}$$

Appendix B. Finite volume discretization for the branching model

$$\frac{du_{1,n_V}}{dt} = \frac{1}{h^2} D_{1,n_V-\frac{1}{2}} (u_{1,n_V-1} - u_{1,n_V}) - \frac{1}{2h} v_{1,n_V-\frac{1}{2}} (u_{1,n_V-1} + u_{1,n_V}) + g_{1,n_V} u_{1,n_V}$$

$$\begin{aligned} \frac{du_{2,n_a}}{dt} &= -\frac{1}{h^2} D_{2,n_a+\frac{1}{2}} (u_{2,n_a} - u_{1,n_a+1}) - \frac{1}{2h} v_{2,n_a+\frac{1}{2}} (u_{2,n_a} + u_{2,n_a+1}) + g_{2,n_a} u_{2,n_a} \\ &\quad - \delta_{12} u_{1,n_a} + \delta_{21} u_{2,n_a} \end{aligned}$$

$$\begin{aligned} \frac{du_{2,i}}{dt} &= \frac{1}{h^2} \left(D_{2,i-\frac{1}{2}} (u_{2,i-1} - u_{2,i}) - D_{2,i+\frac{1}{2}} (u_{2,i} - u_{2,i+1}) \right) \\ &\quad + \frac{1}{2h} \left(v_{2,i-\frac{1}{2}} (u_{2,i-1} + u_{2,i}) - v_{2,i+\frac{1}{2}} (u_{2,i} + u_{2,i+1}) \right) \\ &\quad + g_{2,i} u_{2,i} - \delta_{12} u_{1,i} + \delta_{21} u_{2,i} \end{aligned}$$

for $i \in \{n_a + 1, \dots, n_e\}$

$$\begin{aligned} \frac{du_{2,i}}{dt} &= \frac{1}{h^2} \left(D_{2,i-\frac{1}{2}} (u_{2,i-1} - u_{2,i}) - D_{2,i+\frac{1}{2}} (u_{2,i} - u_{2,i+1}) \right) \\ &\quad + \frac{1}{2h} \left(v_{2,i-\frac{1}{2}} (u_{2,i-1} + u_{2,i}) - v_{2,i+\frac{1}{2}} (u_{2,i} + u_{2,i+1}) \right) + g_{2,i} u_{2,i} \end{aligned}$$

for $i \in \{n_e + 1, n_V - 1\}$

$$\begin{aligned} \frac{du_{2,n_V}}{dt} &= \frac{1}{h^2} D_{2,n_V-\frac{1}{2}} (u_{2,n_V-1} - u_{2,n_V}) \\ &\quad - \frac{1}{2h} v_{2,n_V-\frac{1}{2}} (u_{2,n_V-1} + u_{2,n_V}) + g_{2,n_V} u_{2,n_V}. \end{aligned}$$

Bibliography

- P.-A. Absil, R. Mahony, and R. Sepulchre. *Optimization Algorithms on Matrix Manifolds*. Princeton University Press, 41 William Street, Princeton, New Jersey, 2008.
- R. G. Airapetyan, A. G. Ramm, and A. B. Smirnova. Continuous methods for solving nonlinear ill-posed problems. In A. G. Ramm, P. N. Shivakumar, and A. V. Strauss, editors, *Operator Theory and Applications*, volume 25 of *Fields Institute Communications*, pages 111–138. Amer. Math. Soc., 2000.
- H. Akaike. On the likelihood of a time series model. *J. R. Stat. Soc. Ser. D Stat.*, 27(3/4): 217–235, 1978.
- D. Alvarez, E. H. Vollmann, and U. H. von Andrian. Mechanisms and consequences of dendritic cell migration. *Immunity*, 29(3):325–342, 2008.
- H. Amann. *Linear and Quasilinear Parabolic Problems: Volume I: Abstract Linear Theory*, volume 89 of *Monographs in Mathematics*. Birkhäuser Basel, 1995.
- C. Andres, S. Meyer, O. A. Dina, J. D. Levine, and T. Hucho. Quantitative automated microscopy (QuAM) elucidates growth factor specific signalling in pain sensitization. *Mol. Pain*, 6(98):1–16, 2010.
- C. Andres, J. Hasenauer, H.-S. Ahn, E. K. Joseph, J. Isensee, F. J. Theis, F. Allgöwer, J. D. Levine, S. D. Dib-Hajj, S. G. Waxman, and T. Hucho. Wound healing growth factor, basic FGF, induces Erk1/2 dependent mechanical hyperalgesia. *Pain*, 154(10): 2216–2226, 2013.
- P. Angerer, L. Simon, S. Tritschler, F. A. Wolf, D. Fischer, and F. J. Theis. Single cells make big data: New challenges and opportunities in transcriptomics. *Curr. Opin. Syst. Biol.*, 4:85–91, 2017.
- J. Bachmann, A. Raue, M. Schilling, M. E. Böhm, C. Kreutz, D. Kaschek, H. Busch, N. Gretz, W. D. Lehmann, J. Timmer, and U. Klingmüller. Division of labor by dual

Bibliography

- feedback regulators controls JAK2/STAT5 signaling over broad ligand range. *Mol. Syst. Biol.*, 7(1):516, 2011.
- T. Bäck. *Evolutionary algorithms in theory and practice: evolution strategies, evolutionary programming, genetic algorithms*. Oxford University Press, 1996.
- E. Balsa-Canto, M. Peifer, J. R. Banga, J. Timmer, and C. Fleck. Hybrid optimization method with general switching strategy for parameter estimation. *BMC Syst. Biol.*, 2(26), 2008.
- D. R. Bandura, V. I. Baranov, O. I. Ornatsky, A. Antonov, R. Kinach, X. Lou, S. Pavlov, S. Vorobiev, J. E. Dick, and S. D. Tanner. Mass cytometry: Technique for real time single cell multitarget immunoassay based on inductively coupled plasma time-of-flight mass spectrometry. *Anal. Chem.*, 81(16):6813–6822, 2009.
- J. R. Banga. Optimization in computational systems biology. *BMC Syst. Biol.*, 2(47), 2008.
- H. T. Banks and K. Kunsich. *Estimation techniques for distributed parameter systems*. Systems & Control: Foundations & Applications. Birkhäuser Basel, 1989.
- D. P. Bertsekas. *Nonlinear programming*. Athena Scientific, 2nd edition, 1999.
- H. Bock, T. Carraro, W. Jäeger, S. Körkel, R. Rannacher, and J. Schlöder. *Model Based Parameter Estimation: Theory and Applications*, volume 4 of *Contributions in Mathematical and Computational Sciences*. Springer-Verlag Berlin Heidelberg, 2013.
- B. Bodenmiller, E. R. Zunder, R. Finck, T. J. Chen, E. S. Savig, R. V. Bruggner, E. F. Simonds, S. C. Bendall, K. Sachs, P. O. Krutzik, and G. P. Nolan. Multiplexed mass cytometry profiling of cellular states perturbed by small-molecule regulators. *Nat. Biotechnol.*, 30(9):858–867, 2012.
- R. Boiger, J. Hasenauer, S. Hroß, and B. Kaltenbacher. Integration based profile likelihood calculation for PDE constrained parameter estimation problems. *Inverse Prob.*, 32(12):125009, 2016.
- R. Boiger, A. Fiedler, J. Hasenauer, and B. Kaltenbacher. Continuous analogue to iterative optimization for PDE-constrained inverse problems. *Inverse Probl. Sci. Eng.*, 27(6):710–734, 2019.
- C. Botsaris. A class of methods for unconstrained minimization based on stable numerical integration techniques. *J. Math. Anal. Appl.*, 63(3):729–749, 1978.

- S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, UK, 2004.
- A. A. Brown and M. C. Bartholomew-Biggs. Some effective methods for unconstrained optimization based on the solution of systems of ordinary differential equations. *J. Optim. Theory Appl.*, 62(2):211–224, 1989.
- K. P. Burnham and D. R. Anderson. *Model selection and multimodel inference: A practical information-theoretic approach*. Springer-Verlag New York, 2nd edition, 2002.
- R. H. Byrd, M. E. Hribar, and J. Nocedal. An interior point algorithm for large-scale nonlinear programming. *SIAM J. Optim.*, 9(4):877–900, 1999.
- R. H. Byrd, J. C. Gilbert, and J. Nocedal. A trust region method based on interior point techniques for nonlinear programming. *Math. Program.*, 89(1):149–185, 2000.
- E. P. Carvalho, J. Martínez, J. M. Martínez, and F. Pisnitchenko. On optimization strategies for parameter estimation in models governed by partial differential equations. *Math Comput Simul.*, 114:14–24, 2015.
- J. C. Chambard, R. Lefloch, J. Pouysségur, and P. Lenormand. ERK implication in cell cycle regulation. *Biochim. Biophys. Acta*, 1773(8):1299–1310, 2007.
- J.-S. Chen and R. I. Jennrich. Simple accurate approximation of likelihood profiles. *J. Comput. Graphical Statist.*, 11(3):714–732, 2002.
- O.-T. Chis, J. R. Banga, and E. Balsa-Canto. Structural identifiability of systems biology models: A critical comparison of methods. *PLoS ONE*, 6(11):e27755, 2011.
- H. Cho, K. Ayers, L. de Pills, Y.-H. Kuo, J. Park, A. Radunskaya, and R. C. Rockne. Modelling acute myeloid leukaemia in a continuum of differentiation states. *Lett. Biomath.*, 5(sup1):S69–S98, 2018.
- I.-C. Chou and E. O. Voit. Recent developments in parameter estimation and structure identification of biochemical and genomic systems. *Math. Biosci.*, 219(2):57–83, 2009.
- K.-C. Chou. Applications of graph theory to enzyme kinetics and protein folding kinetics: Steady and non-steady-state systems. *Biophys. Chem.*, 35(1):1–24, 1990.
- A. M. Cook. *Proliferation and lineage potential in fetal thymic epithelial progenitor cells*. PhD thesis, The University of Edinburgh, 2010.

Bibliography

- A. Cornish-Bowden. An automatic method for deriving steady-state rate equations. *Biochem J.*, 165(1):55–59, 1977.
- M. Doumic, A. Marciniak-Czochra, B. Perthame, and J. P. Zubelli. A structured population model of cell differentiation. *SIAM J. Appl. Math.*, 71(6):1918–1940, 2011.
- H.-B. Dürr and C. Ebenbauer. A smooth vector field for saddle point problems. In *Proceedings of the 50th Conference on Decision and Control (CDC 2011)*, pages 4654–4660, Orlando, Florida, USA, 2011.
- J. A. Egea, D. Henriques, T. Cokelaer, A. F. Villaverde, A. MacNamara, D. P. Danciu, J. R. Banga, and J. Saez-Rodriguez. MEIGO: An open-source software suite based on metaheuristics for global optimization in systems biology and bioinformatics. *BMC Bioinformatics*, 15(136), 2014.
- H. W. Engl, M. Hanke, and A. Neubauer. *Regularization of Inverse Problems*, volume 375 of *Mathematics and Its Applications*. Springer Netherlands, 2000.
- D. Faller, U. Klingmüller, and J. Timmer. Simulation methods for optimal experimental design in systems biology. *Simul.*, 79(12):717–725, 2003.
- M. F. Favata, K. Y. Horiuchi, E. J. Manos, A. J. Daulerio, D. A. Stradley, W. S. Feeser, D. E. Van Dyk, W. J. Pitts, R. A. Earl, F. Hobbs, R. A. Copeland, R. L. Magolda, P. A. Scherle, and J. M. Trzaskos. Identification of a novel inhibitor of mitogen-activated protein kinase kinase. *J. Biol. Chem.*, 273(29):18623–18632, 1998.
- E. Feliu and C. Wiuf. Variable elimination in chemical reaction networks with mass-action kinetics. *SIAM J Appl Math*, 72(4):959–981, 2012.
- A. Fiedler, S. Raeth, F. J. Theis, A. Hausser, and J. Hasenauer. Tailored parameter optimization methods for ordinary differential equation models with steady-state constraints. *BMC Syst. Biol.*, 10(1):80, 2016.
- D. S. Fischer, A. K. Fiedler, E. M. Kernfeld, R. M. J. Genga, A. Bastidas-Ponce, M. Bakhti, H. Lickert, J. Hasenauer, R. Maehr, and F. J. Theis. Inferring population dynamics from single-cell RNA-sequencing time series data. *Nat. Biotechnol.*, 37(4):461–468, 2019.
- R. Fritsche-Guenther, F. Witzel, A. Sieber, R. Herr, N. Schmidt, S. Braun, T. Brummer, C. Sers, and N. Blüthgen. Strong negative feedback from Erk to Raf confers robustness to MAPK signalling. *Mol. Syst. Biol.*, 7(1):489, 2011.

- F. Fröhlich, B. Kaltenbacher, F. J. Theis, and J. Hasenauer. Scalable parameter estimation for genome-scale biochemical reaction networks. *PLoS Comput. Biol.*, 13(1):e1005331, 2017a.
- F. Fröhlich, F. J. Theis, J. O. Rädler, and J. Hasenauer. Parameter estimation for dynamical systems with discrete events and logical operations. *Bioinformatics*, 33(7):1049–1056, 2017b.
- F. Fröhlich, T. Kessler, D. Weindl, A. Shadrin, L. Schmiester, H. Hache, A. Muradyan, M. Schütte, J.-H. Lim, M. Heinig, F. J. Theis, H. Lehrach, C. Wierling, B. Lange, and J. Hasenauer. Efficient parameter estimation enables the prediction of drug response using a mechanistic pan-cancer pathway model. *Cell Syst.*, 7(6):567–579.e6, 2018.
- T. Gardner, C. Cantor, and J. Collins. Construction of a genetic toggle switch in *Escherichia coli*. *Nature*, 403(6767):339–342, 2000.
- H. Georgii. *Stochastik: Einführung in die Wahrscheinlichkeitstheorie und Statistik*. De Gruyter Lehrbuch. De Gruyter, 2009.
- L. Haghverdi, F. Buettner, and F. J. Theis. Diffusion maps for high-dimensional single-cell analysis of differentiation data. *Bioinformatics*, 31(18):2989–2998, 2015.
- L. Haghverdi, M. Büttner, F. A. Wolf, F. Buettner, and F. J. Theis. Diffusion pseudotime robustly reconstructs lineage branching. *Nat. Methods*, 13(10):845–848, 2016.
- Á. M. Halász, H.-J. Lai, M. McCabe Pryor, K. Radhakrishnan, and J. S. Edwards. Analytical solution of steady-state equations for chemical reaction networks with bilinear rate laws. *IEEE/ACM Trans. Comput. Biol. Bioinform.*, 10(4):957–969, 2013.
- J. Hasenauer, C. Hasenauer, T. Hucho, and F. J. Theis. ODE constrained mixture modelling: A method for unraveling subpopulation structures and dynamics. *PLoS Comput. Biol.*, 10(7):e1003686, 2014.
- J. Hasenauer, N. Jagiella, S. Hross, and F. J. Theis. Data-driven modelling of biological multi-scale processes. *J. Coupled Syst. Multiscale Dyn.*, 3(2):101–121, 2015.
- T. Hashimoto, D. Gifford, and T. Jaakkola. Learning population-level diffusions with generative RNNs. In M. F. Balcan and K. Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, pages 2417–2426, New York, New York, USA, 2016. PMLR.

Bibliography

- H. Hass, C. Loos, E. Raimúndez-Álvarez, J. Timmer, J. Hasenauer, and C. Kreutz. Benchmark problems for dynamic modeling of intracellular processes. *Bioinformatics*, 35(17):3073–3082, 2019.
- A. C. Hindmarsh, P. N. Brown, K. E. Grant, S. L. Lee, R. Serban, D. E. Shumaker, and C. S. Woodward. SUNDIALS: Suite of Nonlinear and Differential/Algebraic Equation Solvers. *ACM T. Math. Software.*, 31(3):363–396, 2005.
- M. Hinze, R. Pinnau, M. Ulbrich, and S. Ulbrich. *Optimization with PDE Constraints*, volume 23 of *Mathematical Modelling: Theory and Applications*. Springer Netherlands, 2009.
- S. Hock, J. Hasenauer, and F. J. Theis. Modeling of 2D diffusion processes based on microscopy data: Parameter estimation and practical identifiability analysis. *BMC Bioinformatics*, 14(10):S7, 2013.
- H. G. Holzhütter, D. Drasdo, T. Preusser, J. Lippert, and A. M. Henney. The virtual liver: a multidisciplinary, multilevel challenge for systems biology. *Wiley Interdiscip. Rev. Syst. Biol. Med.*, 4(3):221–235, 2012.
- S. Hroß. *Parameter estimation and uncertainty quantification for image based systems biology*. PhD thesis, Technische Universität München, 2016.
- S. Hross and J. Hasenauer. Analysis of CFSE time-series data using division-, age- and label-structured population models. *Bioinformatics*, 32(15):2321–2329, 2016.
- S. Hross, F. J. Theis, M. Sixt, and J. Hasenauer. Mechanistic description of spatial processes using integrative modelling of noise-corrupted imaging data. *J R Soc Interface*, 15(149):20180600, 2018.
- S. Hug, A. Raue, J. Hasenauer, J. Bachmann, U. Klingmüller, J. Timmer, and F. J. Theis. High-dimensional Bayesian parameter estimation: Case study for a model of JAK2/STAT5 signaling. *Math. Biosci.*, 246(2):293–304, 2013.
- W. Hundsdorfer and J. Verwer. *Numerical Solution of Time-Dependent Advection-Diffusion-Reaction Equations*, volume 33 of *Springer Series in Computational Mathematics*. Springer-Verlag Berlin Heidelberg, 2003.
- B. Ingalls. *Mathematical modelling in systems biology: An introduction*. The MIT Press. MIT Press, 2013.

- V. Isakov. *Inverse Problems for Partial Differential Equations*, volume 127 of *Applied Mathematical Sciences*. Springer-Verlag New York, 2nd edition, 2006.
- K. Ito and K. Kunisch. *Lagrange Multiplier Approach to Variational Problems and Applications*, volume 15 of *Advances in Design and Control*. Society for Industrial and Applied Mathematics, 2008.
- B. Kaltenbacher, A. Neubauer, and A. G. Ramm. Convergence rates of the continuous regularized Gauss-Newton method. *J. Inv. Ill-Posed Problems*, 10(3):261–280, 2002.
- J. R. Karr, J. C. Sanghvi, D. N. Macklin, M. V. Gutschow, J. M. Jacobs, B. Bolival Jr, N. Assad-Garcia, J. I. Glass, and M. W. Covert. A whole-cell computational model predicts phenotype from genotype. *Cell*, 150(2):389–401, 2012.
- H. K. Khalil. *Nonlinear Systems*. Prentice Hall, Upper Saddle River, New Jersey, 2nd edition, 1996.
- B. N. Kholodenko. Negative feedback and ultrasensitivity can bring about oscillations in the mitogen-activated protein kinase cascades. *Eur. J. Biochem.*, 267(6):1583–1588, 2000.
- B. N. Kholodenko. Untangling the signalling wires. *Nat. Cell Biol.*, 9(3):247–249, 2007.
- E. L. King and C. Altman. A schematic method of deriving the rate laws for enzyme-catalyzed reactions. *J. Phys. Chem.*, 60(10):1375–1378, 1956.
- H. Kitano. Systems biology: A brief overview. *Science*, 295(5560):1662–1664, 2002.
- A. M. Klein, L. Mazutis, I. Akartuna, N. Tallapragada, A. Veres, V. Li, L. Peshkin, D. A. Weitz, and M. W. Kirschner. Droplet barcoding for single-cell transcriptomics applied to embryonic stem cells. *Cell*, 161(5):1187–1201, 2015.
- E. Klipp, R. Herwig, A. Kowald, C. Wierling, and H. Lehrach. *Systems biology in practice: Concepts, implementation and application*. John Wiley & Sons, Ltd, 2005.
- T. Kose. Solutions of saddle value problems by differential equations. *Econometrica*, 24(1):59–70, 1956.
- C. Kreutz, M. M. Bartolome Rodriguez, T. Maiwald, M. Seidl, H. E. Blum, L. Mohr, and J. Timmer. An error model for protein quantification. *Bioinformatics*, 23(20):2747–2753, 2007.

Bibliography

- F. Lampariello. On the Kolmogorov-Smirnov statistical test for immunofluorescence histogram comparison. *Cytometry Part A*, 39(3):179–188, 2000.
- A. D. Lander, K. K. Gokoffski, F. Y. M. Wan, Q. Nie, and A. L. Calof. Cell lineages and the logic of proliferative control. *PLoS Biol.*, 7(1):e1000015, 2009.
- S. Linnarsson and S. A. Teichmann. Single-cell genomics: coming of age. *Genome Biol.*, 17(1):97, 2016.
- P. M. Loriaux, G. Tesler, and A. Hoffmann. Characterizing the relationship between steady state and response using analytical expressions for the steady states of mass action models. *PLoS Comput. Biol.*, 9(2):e1002901, 2013.
- T. Luzyanina, D. Roose, and G. Bocharov. Distributed parameter identification for label-structured cell population dynamics model using CFSE histogram time-series data. *J. Math. Biol.*, 59(5):581–603, 2009.
- T. Luzyanina, J. Cupovic, B. Ludewig, and G. Bocharov. Mathematical models for CFSE labelled lymphocyte dynamics: Asymmetry and time-lag in division. *J. Math. Biol.*, 69(6–7):1547–1583, 2014.
- M. Marin and A. Öchsner. *Essentials of Partial Differential Equations: With Applications*. Springer International Publishing, 1st edition, 2019.
- W. Q. Meeker and L. A. Escobar. Teaching about approximate confidence regions based on maximum likelihood estimation. *Am. Stat.*, 49(1):48–53, 1995.
- I. Mellman and R. M. Steinman. Dendritic cells: specialized and regulated antigen processing machines. *Cell*, 106(3):255–258, 2001.
- P. Metzger, J. Hasenauer, and F. Allgöwer. Modeling and analysis of division-, age-, and label-structured cell populations. In A. Larjo, S. Schober, M. Farhan, M. Bossert, and O. Yli-Harja, editors, *Proc. of 9th International Workshop on Computational Systems Biology*, pages 55–58, Ulm, Germany, 2012. Tampere International Center for Signal Processing.
- A. N. Michel, L. Hou, and D. Liu. *Stability of Dynamical Systems: On the Role of Monotonic and Non-Monotonic Lyapunov Functions*. Birkhäuser Basel, 2nd edition, 2015.
- C. G. Moles, P. Mendes, and J. R. Banga. Parameter estimation in biochemical pathways: A comparison of global optimization methods. *Genome Res.*, 13(11):2467–2474, 2003.

- E. J. Molinelli, A. Korkut, W. Wang, M. L. Miller, N. P. Gauthier, X. Jing, P. Kaushik, Q. He, G. Mills, D. B. Solit, C. A. Pratilas, M. Weigt, A. Braunstein, A. Pagnani, R. Zecchina, and C. Sander. Perturbation biology: Inferring signaling networks in cellular systems. *PLoS Comput. Biol.*, 9(12):e1003290, 2013.
- J. Müller and C. Kuttler. *Methods and Models in Mathematical Biology: Deterministic and Stochastic Approaches*. Springer-Verlag Berlin Heidelberg, 2015.
- S. A. Murphy and A. W. van der Vaart. On profile likelihood. *J. Am. Stat. Assoc.*, 95(450):449–485, 2000.
- G. D. Nicol and M. R. Vasko. Unraveling the story of NGF-mediated sensitization of nociceptive sensory neurons: ON or OFF the Trks? *Mol. Interv.*, 7(1):26–41, 2007.
- B. F. Nielsen, M. Lysaker, and P. Grøttum. Computing ischemic regions in the heart with the bidomain model – first steps towards validation. *IEEE Trans. Med. Imaging*, 32(6):1085–1096, 2013.
- E. L. O’Dea, D. Barken, R. Q. Peralta, K. T. Tran, S. L. Werner, J. D. Kearns, A. Levchenko, and A. Hoffmann. A homeostatic model of κ b metabolism to control constitutive nf- κ b activity. *Mol. Syst. Biol.*, 3(1):111, 2007.
- E. M. Ozbudak, M. Thattai, H. N. Lim, B. I. Shraiman, and A. van Oudenaarden. Multistability in the lactose utilization network of Escherichia coli. *Nature*, 427(6976):737–740, 2004.
- B. Perthame. *Parabolic Equations in Biology: Growth, reaction, movement and diffusion*. Springer International Publishing, 1 edition, 2015.
- C. Pozo, A. Marín-Sanguino, R. Alves, G. Guillén-Gosálbez, L. Jiménez, and A. Sorribas. Steady-state global optimization of metabolic non-linear dynamic models through recasting into power-law canonical models. *BMC Syst. Biol.*, 5(137), 2011.
- A. Raue, C. Kreutz, T. Maiwald, J. Bachmann, M. Schilling, U. Klingmüller, and J. Timmer. Structural and practical identifiability analysis of partially observed dynamical models by exploiting the profile likelihood. *Bioinformatics*, 25(15):1923–1929, 2009.
- A. Raue, C. Kreutz, F. J. Theis, and J. Timmer. Joining forces of Bayesian and frequentist methodology: A study for inference in the presence of non-identifiability. *Philos. Trans. A Math. Phys. Eng. Sci.*, 371(1984):20110544, 2013a.

Bibliography

- A. Raue, M. Schilling, J. Bachmann, A. Matteson, M. Schelke, D. Kaschek, S. Hug, C. Kreutz, B. D. Harms, F. J. Theis, U. Klingmüller, and J. Timmer. Lessons learned from quantitative dynamical modeling in systems biology. *PLoS ONE*, 8(9):e74335, 2013b.
- A. Raue, B. Steiert, M. Schelker, C. Kreutz, T. Maiwald, H. Hass, J. Vanlier, C. Tönsing, L. Adlung, R. Engesser, W. Mader, T. Heinemann, J. Hasenauer, M. Schilling, T. Höfer, E. Klipp, F. J. Theis, U. Klingmüller, B. Schöberl, and J. Timmer. Data2Dynamics: a modeling environment tailored to parameter estimation in dynamical systems. *Bioinformatics*, 31(21):3558–3560, 2015.
- H. Risken and T. Frank. *The Fokker-Planck equation: Methods of solution and applications*. Springer-Verlag Berlin Heidelberg, 2nd edition, 1996.
- M. Rosenblatt, J. Timmer, and D. Kaschek. Customized steady-state constraints for parameter estimation in non-linear ordinary differential equation models. *Front. Cell. Dev. Biol.*, 4:41, 2016.
- Y. Saeys, S. Van Gassen, and B. N. Lambrecht. Computational flow cytometry: helping to make sense of high-dimensional immunology data. *Nat. Rev. Immunol.*, 16(7):449–462, 2016.
- S. D. M. Santos, P. J. Verveer, and P. I. H. Bastiaens. Growth factor-induced MAPK network topology shapes Erk response determining PC-12 cell fate. *Nat. Cell Biol.*, 9(3):324–330, 2007.
- T. E. Saunders, K. Z. Pan, A. Angel, Y. Guan, J. V. Shah, M. Howard, and F. Chang. Noise reduction in the intracellular pom1p gradient by a dynamic clustering mechanism. *Devel. Cell*, 22(3):558–572, 2012.
- G. Schiebinger, J. Shu, M. Tabaka, B. Cleary, V. Subramanian, A. Solomon, J. Gould, S. Liu, S. Lin, P. Berube, L. Lee, J. Chen, J. Brumbaugh, P. Rigollet, K. Hochedlinger, R. Jaenisch, A. Regev, and E. S. Lander. Optimal-transport analysis of single-cell gene expression identifies developmental trajectories in reprogramming. *Cell*, 176(4):928–943.e22, 2019.
- K. Schumann, T. Lämmermann, M. Bruckner, D. F. Legler, J. Polleux, J. P. Spatz, G. Schuler, R. Förster, M. B. Lutz, L. Sorokin, and M. Sixt. Immobilized chemokine fields and soluble chemokine gradients cooperatively shape migration patterns of dendritic cells. *Immunity*, 32(5):703–713, 2010.

- G. Schwarz. Estimating the dimension of a model. *Ann. Statist.*, 6(2):461–464, 1978.
- A. E. Sgro, D. J. Schwab, J. Noorbakhsh, T. Mestler, P. Mehta, and T. Gregor. From intracellular signaling to population oscillations: bridging size- and time-scales in collective behavior. *Mol. Syst. Biol.*, 11(1):779, 2015.
- O. Shalem, N. E. Sanjana, and F. Zhang. High-throughput functional genomics using CRISPR–Cas9. *Nat. Rev. Genet.*, 16(5):299–311, 2015.
- F. Shiraishi, E. Yoshida, and E. O. Voit. An efficient and very accurate method for calculating steady-state sensitivities in metabolic reaction systems. *IEE/ACM Trans. Comput. Biol. Bioinform.*, 11(6):1077–1086, 2014.
- J. W. Sinko and W. Streifer. A new model for age-size structure of a population. *Ecology*, 48(6):910–918, 1967.
- P. Stapor, F. Fröhlich, and J. Hasenauer. Optimization and profile calculation of ODE models using second order adjoint sensitivity analysis. *Bioinformatics*, 34(13):i151–i159, 2018a.
- P. Stapor, D. Weindl, B. Ballnus, S. Hug, C. Loos, A. Fiedler, S. Krause, S. Hroß, F. Fröhlich, and J. Hasenauer. PESTO: Parameter ESTimation TOolbox. *Bioinformatics*, 34(4):705–707, 2018b.
- K. Tanabe. Continuous Newton-Raphson method for solving an underdetermined system of nonlinear equations. *Nonlinear Anal. Theory Methods Appl.*, 3(4):495–503, 1979.
- K. Tanabe. A geometric method in nonlinear programming. *J. Optim. Theory Appl.*, 30(2):181–210, 1980.
- K. Tanabe. Global analysis of continuous analogues of the Levenberg-Marquardt and Newton-Raphson methods for solving nonlinear equations. *Ann. Inst. Stat. Math.*, 37(1):189–203, 1985.
- A. Tarantola. *Inverse Problem Theory and Methods for Model Parameter Estimation*. Society for Industrial and Applied Mathematics, 2005.
- R. Temam. *Infinite-Dimensional Dynamical Systems in Mechanics and Physics*, volume 68 of *Applied Mathematical Sciences*. Springer-Verlag New York, 2nd edition, 1997.
- C. Trapnell, D. Cacchiarelli, J. Grimsby, P. Pokharel, S. Li, M. Morse, N. J. Lennon, K. J. Livak, T. S. Mikkelsen, and J. L. Rinn. The dynamics and regulators of cell fate

Bibliography

- decisions are revealed by pseudotemporal ordering of single cells. *Nat. Biotechnol.*, 32(4):381–386, 2014.
- J. Vanlier, C. A. Tiemann, P. A. J. Hilbers, and N. A. W. van Riel. An integrated strategy for prediction uncertainty analysis. *Bioinformatics*, 28(8):1130–1135, 2012.
- A. I. F. Vaz and L. N. Vicente. A particle swarm pattern search method for bound constrained global optimization. *J. Glob. Optim.*, 39(2):197–219, 2007.
- D. J. Venzon and S. H. Moolgavkar. A method for computing profile-likelihood-based confidence intervals. *J. R. Stat. Soc. Ser. C. Appl. Stat.*, 37(1):87–94, 1988.
- A. F. Villaverde, D. Henriques, K. Smallbone, S. Bongard, J. Schmid, D. Cicin-Sain, A. Crombach, J. Saez-Rodriguez, K. Mauch, E. Balsa-Canto, P. Mendes, J. Jaeger, and J. R. Banga. BioPreDyn-bench: A suite of benchmark problems for dynamic modelling in systems biology. *BMC Syst. Biol.*, 9(8), 2015.
- A. F. Villaverde, F. Fröhlich, D. Weindl, J. Hasenauer, and J. R. Banga. Benchmarking optimization methods for parameter estimation in large kinetic models. *Bioinformatics*, 35(5):830–838, 2018.
- A. Wächter and L. T. Biegler. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Math. Program.*, 106(1):25–57, 2006.
- W. Walter. *Gewöhnliche Differentialgleichungen: Eine Einführung*. Springer-Verlag Berlin Heidelberg, 5th edition, 1993.
- R. A. Waltz, J. L. Morales, J. Nocedal, and D. Orban. An interior algorithm for nonlinear optimization that combines line search and trust region steps. *Math. Program.*, 107(3):391–408, 2006.
- S. Wang, X. Q. Yang, and K. L. Teo. A unified gradient flow approach to constrained nonlinear optimization problems. *Comput. Optim. Appl.*, 25(1):251–268, 2003.
- J. V. Watson. Proof without prejudice revisited: immunofluorescence histogram analysis using cumulative frequency subtraction plus ratio analysis of means. *Cytometry*, 43(1):55–68, 2001.
- M. Weber, R. Hauschild, J. Schwarz, C. Moussion, I. de Vries, D. F. Legler, S. A. Luther, T. Bollenbach, and M. Sixt. Interstitial dendritic cell guidance by haptotactic chemokine gradients. *Science*, 339(6117):328–332, 2013.

- C. Weinreb, S. Wolock, B. K. Tusi, M. Socolovsky, and A. M. Klein. Fundamental limits on dynamic inference from single-cell snapshots. *Proc. Natl. Acad. Sci. USA*, 115(10):E2467–E2476, 2018.
- S. M. Wilhelm, L. Adnane, P. Newell, A. Villanueva, J. M. Llovet, and M. Lynch. Preclinical overview of sorafenib, a multikinase inhibitor that targets both Raf and VEGF and PDGF receptor tyrosine kinase signaling. *Mol. Cancer Ther.*, 7(10):3129–3140, 2008.
- D. H. Wolpert and W. G. Macready. No free lunch theorems for optimization. *IEEE Trans. Evol. Comput.*, 1(1):67–82, 1997.
- G. Xu. Steady-state optimization of biochemical systems through geometric programming. *European J. Oper. Res.*, 225(1):12–20, 2013.
- X. Xun, J. Cao, B. Mallick, R. J. Carroll, and A. Maity. Parameter Estimation of Partial Differential Equation Models. *J. Am. Stat. Assoc.*, 108(503):1009–1020, 2013.
- X.-S. Yang. *Nature-inspired metaheuristic algorithms*. Luniver Press, Bristol, UK, 2nd edition, 2010.
- M. A. Yui and E. V. Rothenberg. Developmental gene networks: a triathlon on the course to T cell identity. *Nat. Rev. Immunol.*, 14(8):529–545, 2014.
- C. Zechner, J. Ruess, P. Krenn, S. Pelet, M. Peter, J. Lygeros, and H. Koeppl. Moment-based inference predicts bimodality in transient gene expression. *Proc. Natl. Acad. Sci. USA*, 109(21):8340–8345, 2012.
- E. Zeidler. *Nonlinear Functional Analysis and its Applications II/B: Nonlinear Monotone Operators*. Springer-Verlag New York, 1990.
- W. Zhang and H. T. Liu. MAPK signal pathways in the regulation of cell proliferation in mammalian cells. *Cell Res.*, 12(1):9–18, 2002.