Available online at www.sciencedirect.com

ScienceDirect

Procedia Technology 15 (2014) 474 – 483

Procedia Technology

2nd International Conference on System-Integrated Intelligence: Challenges for Product and Production Engineering

First steps towards an intelligent laser welding architecture using deep neural networks and reinforcement learning

Johannes Günther[a]*, Patrick M. Pilarski[b], Gerhard Helfrich[a], Hao Shen[a], Klaus Diepold[a]

[a]*Technische Universität München, Arcisstraße 21, Munich, 80333, Germany*
[b]*Dept. of Computing Science, University of Alberta, Edmonton, AB T6G 2E8, Canada*

Abstract

To address control difficulties in laser welding, we propose the idea of a self-learning and self-improving laser welding system that combines three modern machine learning techniques. We first show the ability of a deep neural network to extract meaningful, low-dimensional features from high-dimensional laser-welding camera data. These features are then used by a temporal-difference learning algorithm to predict and anticipate important aspects of the system's sensor data. The third part of our proposed architecture suggests using these features and predictions to learn to deliver situation-appropriate welding power; preliminary control results are demonstrated using a laser-welding simulator. The intelligent laser-welding architecture introduced in this work has the capacity to improve its performance without further human assistance and therefore addresses key requirements of modern industry. To our knowledge, it is the first demonstrated combination of deep learning and Nexting with general value functions and also the first usage of deep learning for laser welding specifically and production engineering in general. This work also provides a unique example of how predictions can be explicitly learned using reinforcement learning to support laser welding. We believe that it would be straightforward to adapt our approach to other production engineering applications.

© 2014 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY-NC-ND license (http://creativecommons.org/licenses/by-nc-nd/3.0/).
Peer-review under responsibility of the Organizing Committee of ICIAME 2014.

Keywords: Deep learning; reinforcement learning; prediction; control; laser welding

* Corresponding author. Tel.: +49 89 289 23631; fax: +49 289 23600.
E-mail address: johannes.guenther@tum.de.

2212-0173 © 2014 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY-NC-ND license (http://creativecommons.org/licenses/by-nc-nd/3.0/).
Peer-review under responsibility of the Organizing Committee of ICIAME 2014.
doi:10.1016/j.protcy.2014.09.007

1. Introduction

Laser welding is a precise and fast welding technique that sees widespread use in industrial welding systems [1]. Unfortunately, laser welding is a complex process that is often hard to control [2]. To address control difficulties, recent research has demonstrated cognitive laser welding systems that perform well on a defined work piece after setup [3]. Nevertheless, cognitive control is still in an early stage of development [4], and to fulfill the requirements of modern industry, systems must have the flexibility to deal with changing conditions without the need for demanding and time-intensive manual setup [5].

To address the need for both rapid setup times and welding system flexibility, we propose the idea of a self-learning and self-improving laser-welding system that would be able to perform well under changing circumstances. As a classical approach by a model is not feasible due to the dynamics and uncertainty inherent to the process, we suggest applying machine learning techniques. Our proposed approach brings together a selection of modern machine learning techniques, including deep-learning neural networks for generating state representations and state-of-the-art reinforcement learning prediction and control algorithms. These algorithms empower the system to leverage important aspects of intelligence during welding, namely perception, prediction, and interaction.

As laser-welding systems' sensor signals are multidimensional and multimodal, it is often not realistic to use them directly as an input for real-time control learning algorithms. Building on established ideas in dimensionality reduction, we therefore use a representation-learning (perception) algorithm to transform the raw sensor data into a low-dimensional and more invariant representation of the system's state. The system learns to abstract its inputs. In particular, a technique that has shown its capability to produce the lowest classification error for various problems when used for feature extraction is deep learning [6]. Furthermore, deep autoencoders have been shown to successfully compete with state-of-the art feature extraction techniques (e.g., PCA, LDA) [7] and improved [8] or directly learned [9] policies for high-dimensional image data in reinforcement learning. Stacked denoising autoencoders have shown the capability of achieving a general representation, which leads to more robustness against varying data and overfitting [10].

A very common problem in industry is the inability to directly measure process quality. There are several approaches to this issue, e.g. system models, envelop curves or look up tables. But these techniques are restricted either in applicability (a priori model), accuracy (envelop curves) or scalability (look up tables). They are also limited in their capability to adapt to changes. To deal with these issues, we included predictions about process quality and state as an important part of intelligence [11] in our architecture; importantly, we suggest that these predictions should be able to be learned and adapted during the ongoing operation of the system. To date, prediction learning has been dominated by linear models that are difficult to apply to nonlinear and time-varying problems [12]. These problems have been overcome by recent research using the temporal-difference (TD) reinforcement learning approach [13]. New techniques have extended classical TD to allow generalized online predictions [14]. We include these predictions into our proposed system using a temporally extended prediction approach called Nexting [15] with general value functions, an approach that is capable of learning and making real-time predictions at multiple timescales.

There exist a number of different controllers for industrial applications, e.g. PID-controllers, adaptive controllers and fuzzy controllers. Given a correct and accessible quality measurement, it would be easy to implement these techniques for laser welding. But all these approaches need a time-consuming and human assisted set up process and do not work well for changing conditions. To enable our architecture to provide a high-quality welding seam on its own, it is necessary to have a controller that can learn from experience and improve its own performance. Therefore we suggest a machine learning algorithm, namely an actor-critic reinforcement learning (ACRL) algorithm [16]. This type of algorithm consists of two parts: an actor and a critic. The actor takes actions according to a learned policy while the critic evaluates these actions. The actor-critic algorithm has several characteristics that are useful for our specific control problem. As ACRL algorithms are parameter based, their computational cost increments linearly and they can be updated within milliseconds. Due to the fact that experience—from which the algorithm already had learned—does need not be stored, the memory requirements do not increase over time [17]. By using function approximation they also scale well to real world problems; this has been shown in various applications [18], [19], [20], [21].

Our proposed architecture for integrating representation, prediction and control in laser welding therefore promises to address key industry needs relating to both the calibration and optimization of diverse welding processes. It is described in the remainder of this manuscript as follows. Section 2 describes the laser welding system and the monitoring, as well as how the algorithms will work together in the proposed architecture. Section 3 focuses on deep learning and how features are generated via deep autoencoders from the existing sensor input. These features are the input for the reinforcement learning algorithms, explained and evaluated in Section 4. The results are discussed in Section 5 and followed by concluding remarks in Section 6.

2. Laser welding and the proposed architecture

2.1. The laser welding process and monitoring

Although laser welding is quite common in industrial applications, it is still necessary to closely and consistently monitor and control the process [22]. Despite the environmental uncertainties that the process is exposed to, like temperature, changing humidity or changes in the welding gas quality, there are also uncertainties caused by the material. These include, but are not limited to, changes in the chemical compounding, and the thickness and contamination of the surface.

In our setting, process monitoring is done by a camera-based system and photodiodes, which is a common setting in laser welding applications [23]. As the keyhole, which is the area where the laser hits the material, oscillates with a typical frequency of 500 Hz [24], all sensors have to sample with at least twice this frequency. This can be considered as a benchmark real-time capability for the process. The camera can sample at rates of up to 1500 Hz. It provides important information about geometrical parameters of the observed keyhole [25] with a resolution of 144x176 pixels. Additionally, the process is observed by three photodiodes, sampling at 40 kHz and corresponding to different wavelengths. The first diode observes the process temperature at the wavelength between 1100nm-1800nm. The second observes the plasma radiation at a wavelength of 400nm-600nm. The last diode records the laser back reflection at about 1050nm-1080nm. For initial process control results, the control algorithm is tested on a preliminary simulator [26], which provides the welding seam depth based on the welding seam width.

2.2. Architecture

Laser welding is a dynamic process with high uncertainty and therefore it is not feasible to build a precise model of the process, which would be the classical control approach. We therefore propose a machine learning approach. Our suggested architecture combines deep neural networks (DNN) [7] with reinforcement learning algorithms [27].

Figure 1 shows the architecture, which consists of three parts: representation, process knowledge (prediction), and process control. In the first part, deep learning of representations, the monitored sensor data is processed and transformed into informative features which are lower in dimension to ensure real-time capability and robustness. By doing so, the system is able to detect its current state only by the provided sensor data and is therefore more invariant to environmental changes. These features are used in the second part—prediction—to build up knowledge about the process. By using temporally extended predictions, the system is able to evaluate its current performance and predict how its actions might impact its performance in the future. The features from the representation and the knowledge from the second part are combined in the third part—process control—to control the system in terms of the laser power applied to the welding surface.

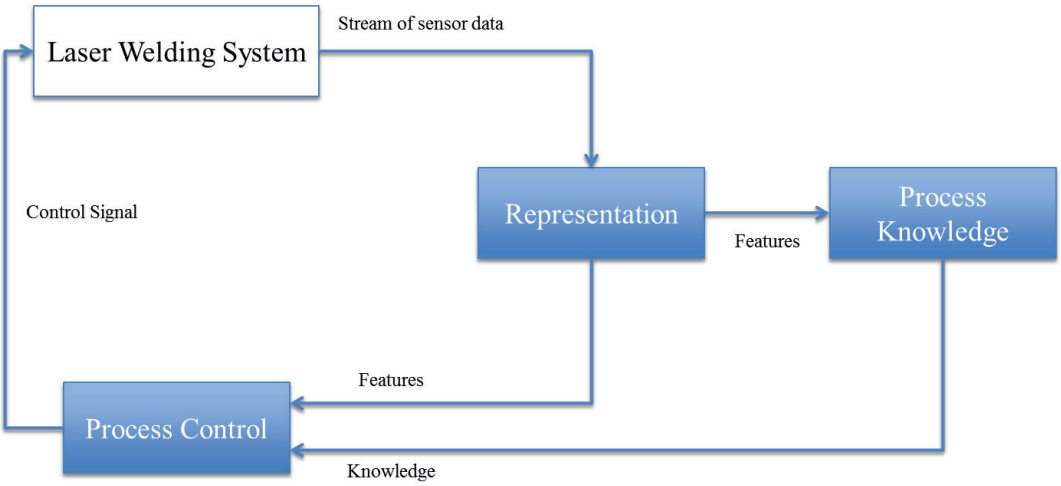

Fig. 1. Proposed architecture for intelligent laser welding. The architecture consists of three elements: representation, process knowledge (prediction), and process control. The first element extracts meaningful low-dimensional features from the sensor data. The second element then uses these features to learn about the process and build up knowledge. This knowledge together with the features is used to control the process in the third element. All machine learning elements are denoted by blue boxes.

3. Deep Learning

Our representation approach employs techniques from deep learning. Deep learning is an advancement of neural networks, and in particular, the autoencoder was introduced be Hinton and Rumelhart as a nonlinear generalization of the PCA [7]. It consists of two parts, the encoder and the decoder. In the encoder, the sensor data is fed into the network via the input layer. The first hidden layer is bigger, i.e. the number of units (artificial neurons) is bigger than in the input layer, to allow a more general representation of the input. The subsequent layers consistently decrease in size to compress the information. In the smallest layer, called the bottleneck, the data is encoded into a low dimensional representation, which can be used as features. The following layers expand in size and reconstruct the original input data from the features in the bottleneck. By doing so, it is ensured that the features contain key and invariant information regarding the original data. The exact parameters for the autoencoder used in our present study can be found in Table 1.

Table 1. Deep autoencoder parameters.

Parameter	Value
Neurons per layer (encoder and bottleneck)	1024-2048-1024-512-256-16
Decay learning rate	0.0005 – 0.00025
Momentum rate [28]	0.005 – 1
L1 weight decay factor	0.00005
Mini-batch size [28]	Pretraining: 100, fine-tuning: 200
Additive Gaussian noise in layer 1 [10]	$N(0, 0.5)$
Additive Gaussian noise in layer 2-4	$N(0, 0.1\ \sigma^k)$
Corruption rate	50%

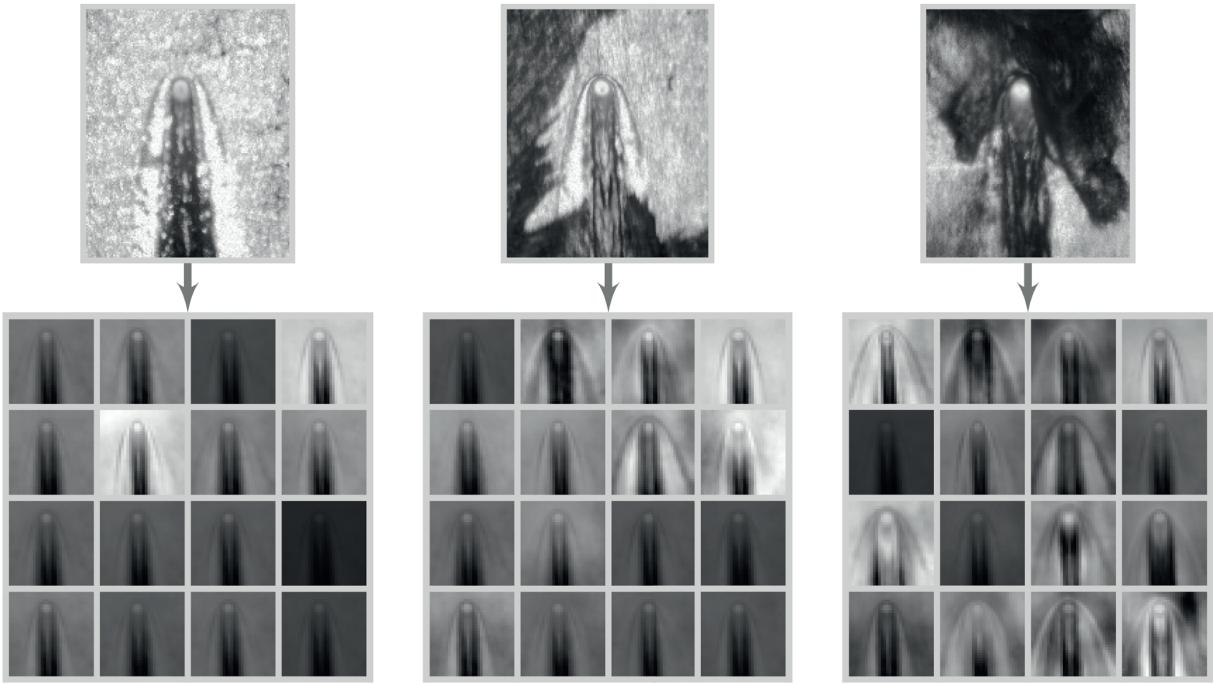

Fig. 2. DNN feature visualization and original images. For each input image 16 features-images are reconstructed, which show the specialization and activation for the image. It can be seen that not only the level of activation but also the specialization for each feature is dependent on the input image. Different features can specialize on different aspects of the input image.

The deep autoencoder achieves a mean reconstruction error of 16.6% from 16-dimensional features on 8 different test data sets. Compared to the PCA, which has a mean reconstruction error of 15.5%, the deep autoencoder is a competitive alternative. The deep autoencoder features also yielded a lower classification error when used as input for two Support Vector Machine (SVM) classifiers (see table 2).

Table 2. Classification error for PCA and deep autoencoder extracted features.

Classification error	PCA	Deep Autoencoder
SVM Sigmoidal Kernel (16 Features)	8.24% ± 0.06%	8.07% ± 0.06%
SVM RBF Kernel (16 Features)	1.79% ± 0.19%	0.81% ± 0.11%

To give insights into the learned DNN, we visualize its features in a specific way: All, except one, features are set to zero. The resulting image is scaled by the value of the remaining, active feature. Using this technique, we generate an image, which shows the specialization and activity of the particular feature as bright parts. This is repeated for every feature. The output visualization for this process can be found in Figure 2.

4. Reinforcement Learning

Reinforcement learning (RL) is a branch of machine learning and artificial intelligence that defines a learning goal but does not explicitly rely on labeled training examples during learning—it learns through trial and error. As such, RL does not necessarily need the designer to provide external domain information for each learning scenario. As RL can consider the whole problem without explicitly dividing and optimizing sub-problems, it is well suited for changing environments and varying conditions [27].

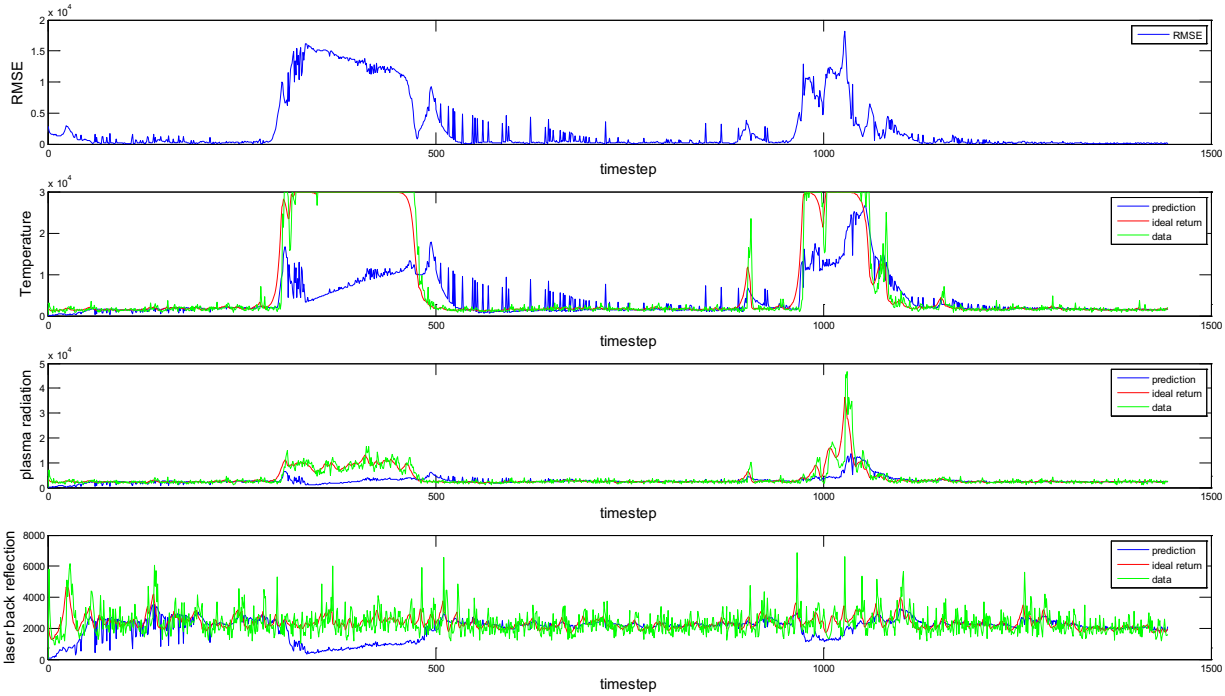

Fig. 3. Welding process with an oil contaminated surface. The process runs for 1444 images (three seconds). The topmost graph shows the cumulated RMSE over time. The other graphs show the performance for each of the photodiodes. The first photodiode records the temperature, the second the plasma and the third the laser back reflex. The green lines represent the actual photodiode data, while the red lines are the ideal prediction and the blue line is the prediction made by the Nexting algorithm.

Reinforcement learning scenarios consist of a learning agent and an environment. The agent affects the environment, using actions, according to a policy π, from a set of possible actions A. The environment responds to every action by changing its state, out of a state space set S, and provides a reward r, which is used by the agent to determine the expected future outcome of the chosen action. The agent maps states to rewards via a learned value function V, and uses this value function to select actions or modify its policy (as in ACRL).

4.1. Prediction

The online predictions about welding process sensors are learned and made with a specific RL algorithmic approach known as Nexting [15], which is a term for making a prediction about what will happen next. This algorithm makes use of temporal-difference learning [27] with linear function approximation and general value functions. It is capable of making predictions at multiple time scales [15].

The algorithm determines the value (prediction) for each state by calculating the scalar product of a state feature vector and a learned weight vector. The temporal-difference error δ, which is the difference between the expected state value and the actual state value, is calculated by adding the sensor value and the estimations for the successor state and subtracting the current state. δ is used to adjust the weights so that the prediction will match the expected sensor values. For the experiment we used a data set containing a three-second process of laser welding. For the process, camera data and photodiode data is available. The camera data is processed by the DNN, tile coded, which is a common linear function approximation technique [27], and used as the prediction input. The photodiode data is the prediction goal, which means the algorithm predicts the future expected photodiode data, based on the camera data five steps in advance (T = 5). As we do this, we verify the meaningfulness of the extracted features and prove the algorithm's capability to improve its predictions.

To have a benchmark, the ideal prediction is computed offline and marked as a red line in Figure 3. As the ideal prediction takes the complete data sequence into consideration, it will anticipate the actual data behavior. The blue line shows the prediction made by the online algorithm. If the algorithm is able to learn about the process and performs well, the online prediction will closely match the ideal prediction.

For the process we used the following learning parameters: learning rate $\alpha = 0.2/m$, discount rate $\gamma = 0.8$, eligibility trace decay parameter $\lambda = 0.7$, where m is the number of active tiles in the tile coding approximation [27].

In the process, the welding seam is contaminated with oil at two different spots. As the laser hits the oil it starts burning and therefore impairs the sensor data. This can be seen in Figure 3 by the spikes in the photodiode data. As this behavior is new to the algorithm the prediction error increases.

For the first photodiode the algorithm struggles to match the ideal prediction when the failure occurs for the first time. But, while running, the algorithm learns more about the system and is able to predict significantly better when the failure occurs the second time. A similar performance improvement can be seen for the other photodiodes. Additionally, the root mean square error (RMSE) is significantly smaller for the second error. To verify the relevance of the predictions, the normalized mean squared error (NMSE) was computed for all three photodiodes. This measure gives the percent of the variance that is not explained by the prediction [15]. For the process, shown in Figure 3, it is: NMSE1 = 0.25, NMSE2 = 0.6, NMSE3 = 0.98.

It can be seen that the algorithm performs best for the first photodiode. This can be explained by the fact that the influence of oil on the work piece on the temperature diode is the biggest, as the burning oil is within its observation frequency. Therefore the correlation between the camera data and the temperature photodiode data is especially strong which results in better predictions.

4.2. Control

To address the continuous state and action space in the simulator, we selected a continuous ACRL method that used a one dimensional action space. The algorithm can be found in Pilarski et al. [17]. In our setting, it provides a value for the applied laser power based on its input, which is a double precision number derived from the simulator. The actor follows a random control policy $\pi(a|s)$, sampled from a Gaussian distribution with the mean μ and the standard deviation σ, which are linear combinations of a corresponding learned weight vector and the feature vector. The standard deviation is a measure for the exploration of the algorithm. If σ is big, the algorithm is more likely to explore different actions. As exploration steps usually result in a poor performance, the standard deviation has to decrease once the right mean has been found to achieve a stable, good performance. At the same time the critic uses the experience to evaluate the algorithm's performance. For every sample $(s_t, a_t, r_{t+1}, s_{t+1})$ the critic calculates the temporal-difference error δ and uses it to update the value function V as well as the weights for calculating μ and σ. For our experiments we used the following learning parameters: learning rate $\alpha_v = 0.1/m$, $\alpha_w = 0.1/m$, discount rate $\gamma = 0.99$, trace decay parameters $\lambda_a = 0.3$, $\lambda_c = 0.3$ and sigma start $\sigma_c = 1$. All other parameters are initialized to zero. Reward is given based on the distance between the actual welding depth and the desired welding depth. To keep the reward bounded and prevent large TD-errors, which might lead to divergence, it is designed as a sigmoid function ranging between 0 and -0.5. The algorithm is optimistically initialized, which means all states have the same good value in the beginning, i.e., zero. The algorithm runs on the simulator, introduced in Section 2.1. 30 independent runs were performed and averaged. None of the trials diverged.

Figure 4 shows several actor-critic parameters over the learning process. As the algorithm draws its actions, i.e. the applied laser power, from a normal distribution, it takes random actions in the beginning of the learning process. According to the reward it gets for its chosen actions, it shifts μ towards the correct action and decreases σ to take less exploration steps. This results in constant actions and therefore a higher reward. From iteration 16,000, which corresponds to a learning time of 32 seconds, the algorithm performs with a precision of 0.5 mm. It improves to 0.1 after 25,000 iterations (50 seconds). The whole learning process corresponds to 93 seconds in real time.

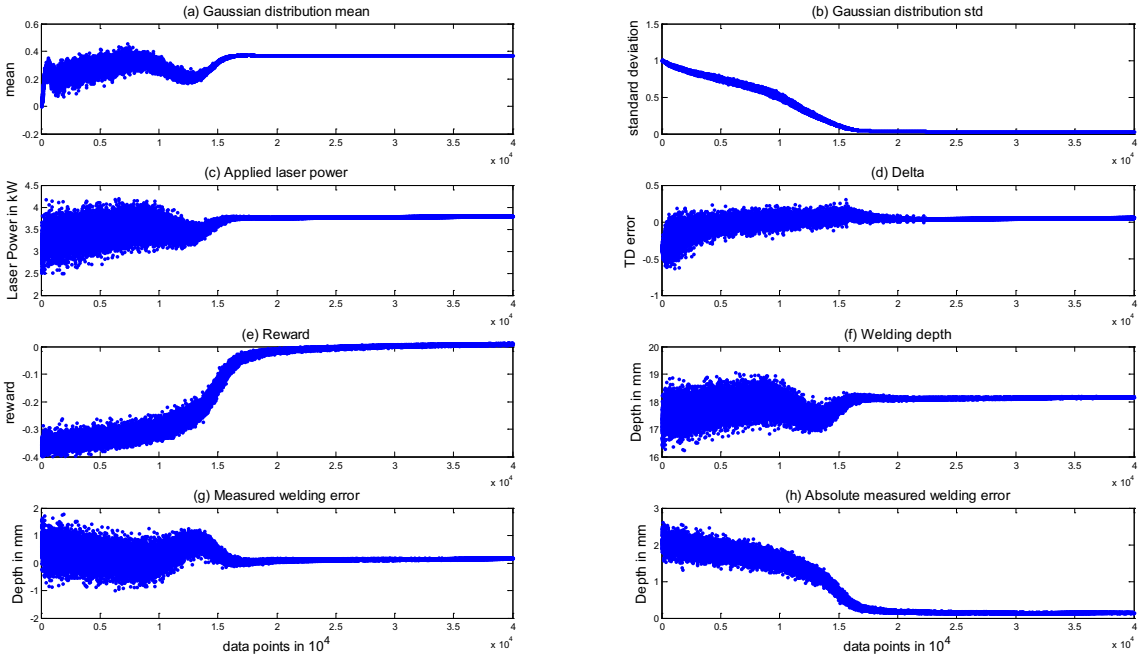

Fig. 4. Average learning performance for the continuous actor-critic algorithm over 30 independent runs. The two topmost plots (a), (b) show the parameters for the Gaussian distribution, shown in plot (c). Plot (d) and (e) give insights, how the algorithm is evaluating its performance, using the temporal difference error und the reward. The plots (f), (g) and (h) visualize the welding performance by showing the welding depth or the distance to the desired welding depth.

5. Discussion

In this paper we proposed a new architecture that should be able to learn and control a wide class of specific laser-welding problems. Although the proposed architecture is flexible enough to deal with varying conditions in the process, the representation and knowledge framework in the present work has been specialized to deal with one specific laser-welding process of interest. Our architecture is capable of being extended as new process-related needs are identified, and we expect the present results will transfer well to different settings.

In Section 3 it was shown that DNN are able to extract meaningful features in a laser-welding setting, which can be used to represent the current process state to other knowledge and control processes. Tests with a SVM classifier showed that the features generated by the DNN perform better than PCA features.

The knowledge process has shown its capability to learn about the welding process's behavior and to predict sensor signals as they extend into the near future. The experiment shown in Figure 3 showed the capability of our algorithmic approach to learn to predict related, temporally extended photodiode data from the features generated by the DNN. The behavior of the learned predictions further validates the meaningfulness of the DNN-generated features. When this approach is used in a real laser-welding environment, it will be able to make use of much larger data sets. We strongly believe that this will result in significantly better and more adaptable predictions. It is also important to note that adding additional sensor dimensions will not necessarily increase the computational complexity, as the state is tile-coded and hashed into a vector whose size is pre-selected. However, not increasing the vector size will make collisions in the hash table more likely and might result in a worse performance (via inappropriate generalization).

The control algorithm described in Section 4.2 consistently converged to the correct solution in a short learning time. It was also fast enough from the computational point of view. The whole simulation takes 66.3 seconds, whereas 52.6 seconds are used to simulate the laser welding system. Therefore, a single control iteration takes

0.34ms, which is within the process requirements for real time capability (1ms). The implementation of the simulator is applicable under the assumption that the representation algorithm is able to provide meaningful features about the systems state and that the prediction process can accurately predict the current performance. Additional work in the industrial setting is needed to demonstrate a fully integrated framework wherein learned DNN features and Nexting predictions are used directly as input to the ACRL control learner during welding.

6. Conclusion

In this paper we have shown a new possible architecture for laser welding. Our proposed system includes methods to observe a process, build up knowledge about it, and then find ways to control it—as such, it can be viewed as a simple form of industrially motivated intelligence. The system has the capacity to improve its own performance, due to the way that it optimizes the process in terms of goals rather than in terms of mechanisms. It therefore promises to addresses key requirements of modern industry, in that our architecture combines fast learning with the capability to work well under changing circumstances.

In this work we described a possible combination of recent reinforcement learning and deep learning algorithms, and provided insights into the impact this combination may have on laser-welding technology. To our knowledge, this is the first demonstrated combination of deep learning with Nexting and general value functions, and one of only a very small number of papers describing the combination of reinforcement learning and deep learning systems, e.g. [8], [9]. Additionally, it is the first demonstrated use of deep learning in laser welding and industrial production processes. This study is also unique in its use of reinforcement learning to acquire generalized predictions for use as inputs to a laser welding system. This makes the present work an important incremental contribution to not only industrial process engineering, but also to the study of intelligent systems and machine intelligence. However, we do not wish to constrain our approach to the exact learning methods deployed in the present work; another important contribution of this paper is to suggest that the full integration of representation, prediction, and control learning into a single framework holds great promise for laser welding and other production engineering domains. We strongly believe that our approach is transferable to a broader range of industrial applications.

Acknowledgements

The project CCLW is performed within the framework of the European funding program Eurostars and is funded by the Federal Ministry of Education and Research. We also would like to thank all project partners, namely the Precitec GmbH & Co. KG, Germany and IREPA LASER, Illkirch, France, who provided the data. The authors also acknowledge support from the Alberta Innovates Centre for Machine Learning and Alberta Innovates Technology Futures (Canada) and thank Richard Sutton and Joseph Modayil for a number of helpful conversations.

References

[1] Jaeger M, Humbert S, Hamprecht FA. Sputter tracking for the automatic monitoring of industrial laser-welding processes. IEEE Transactions on Industrial Electronics 2008; 55: 2177–2184.
[2] Alippi C, Braione P, Piuri V, Scotti F. A methodological approach to multisensor classification for innovative laser material processing units. IMTC 2001; 3: 1762–1767.
[3] Schroth G, Stork I, Wersborg G, Diepold K. A cognitive system for autonomous robotic welding. IEEE/RSJ IROS 2009; 3148–3153.
[4] Haykin S, Fatemi M, Setoodeh P, Xue Y. Cognitive Control. IEEE 2012; 100: 3156–3169.
[5] Zipkin P. The limits of mass customization. Harvard Business Review 1997; 75: 91–101.
[6] Krizhevsky A, Sutskever I, Hinton GE. Imagenet classification with deep convolutional neural networks. Advances in Neural Information Processing Systems 2012; 25: 1106–1114.
[7] Hinton GE, Salakhutdinov RR. Reducing the dimensionality of data with neural networks. Science 2006; 313: 5786: 504–507.
[8] Lange S, Riedmiller M, Voigtlander A. Autonomous reinforcement learning on raw visual input data in a real world application. IJCNN 2012: 1–8.
[9] Mnih V, Kavukcuoglu K, Silver D, Graves A, Antonoglou I, Wierstra D, Riedmiller M. Playing atari with deep reinforcement learning, NIPS Workshop on Deep Learning and Unsupervised Feature Learning 2013.

[10] Vincent P, Larochelle H, Lajoie I, Bengio Y, Manzagol PA. Stacked denoising autoencoders: Learning useful representations in a deep network with a local denoising criterion. Journal of Machine Learning Research 2010; 11: 3371–3408.
[11] Clark A. Whatever next? Predictive brains, situated agents, and the future of cognitive science. Behavioral and Brain Sciences 2013; 36: 03: 181–204.
[12] Liu F, Quek C, Ng GS. Neural network model for time series prediction by reinforcement learning. IEEE IJCNN 2005; 2: 809–814.
[13] Sutton RS, Modayil J, Delp M, T. Degris, Pilarski PM, White A, Precup D. Horde: a scalable real-time architecture for learning knowledge from unsupervised sensorimotor interaction. AAMAS 2011: 761–768.
[14] Modayil J, White A, Pilarski PM, Sutton RS. Acquiring a broad range of empirical knowledge in real time by temporal-difference learning. IEEE SMC 2012: 1903–1910.
[15] Modayil J, White A, Sutton RS. Multi-timescale nexting in a reinforcement learning robot. Adaptive Behavior 2014; 22:2: 146-160.
[16] Degris T, Pilarski PM, Sutton RS. Model-Free reinforcement learning with continuous action in practice. ACC 2012: 2177–2182.
[17 Pilarski PM, Dawson MR, Degris T, Fahimi F, Carey JP, Sutton RS. Online human training of a myoelectric prosthesis controller via actor-critic reinforcement learning. IEEE ICORR 2011: 1–7.
[18] Kimura H, Yamashita T, Kobayashi S. Reinforcement learning of walking behavior for a four-legged robot. CDC 2001; 1: 411–416.
[19] Thomas P, Branicky M, van den Bogert A, Jagodnik K. Application of the actor-critic architecture to functional electrical stimulation control of a human arm. IAAI 2009: 165–172.
[20] Peters J, Schaal S. Reinforcement learning of motor skills with policy gradients. Neural Networks 2008; 21: 4: 682–697.
[21] Pilarski PM, Dick TB, Sutton RS. Real-time prediction learning for the simultaneous actuation of multiple prosthetic joints. IEEE ICORR 2013: 1–8.
[22] Ancona A, Spagnolo V, Lugarà PM, Ferrara M. Optical sensor for real-time monitoring of CO2 laser welding process. Applied Optics 2001; 40: 33: 6019–6025.
[23] Shao J, Yan Y. Review of techniques for on-line monitoring and inspection of laser welding. Journal of Physics 2005; 15: 1: 101.
[24] Kroos J, Gratzke U, Vicanek M, Simon G. Dynamic behaviour of the keyhole in laser welding. Journal of Physics D: Applied Physics 1993; 26: 3: 481.
[25] Beersiek J. A CMOS camera as tool for process analysis not only for laser beam welding. ICALEO 2001; 92: 1185-1193
[26] Min-ying H, Shi-Junwei, Li-xin C, Yanqiu X. The research of welding parameters on weld shape in the laser deep penetration welding. MACE 2010: 3734–3737.
[27] Sutton RS, Barto AG. Reinforcement Learning: An Introduction. MIT Press; 1998.
[28] Bengio Y. Practical recommendations for gradient-based training of deep architectures. Neural Networks 2012: 437-478.