Mechatronics 34 (2016) 1–11

Contents lists available at ScienceDirect

Mechatronics

journal homepage: www.elsevier.com/locate/mechatronics

Intelligent laser welding through representation, prediction, and control learning: An architecture with deep neural networks and reinforcement learning

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

[a] Dept. Electrical and Computer Engineering, Technische Universität München, Munich 80290, Germany
[b] Dept. of Medicine, University of Alberta, Edmonton, AB T6G 2E1, Canada

ARTICLE INFO

Article history:
Received 15 October 2014
Revised 23 July 2015
Accepted 4 September 2015
Available online 1 October 2015

Keywords:
Deep learning
Reinforcement learning
Prediction
Control
Laser welding

ABSTRACT

Laser welding is a widely used but complex industrial process. In this work, we propose the use of an integrated machine intelligence architecture to help address the significant control difficulties that prevent laser welding from seeing its full potential in process engineering and production. This architecture combines three contemporary machine learning techniques to allow a laser welding controller to learn and improve in a self-directed manner. As a first contribution of this work, we show how a deep, auto-encoding neural network is capable of extracting salient, low-dimensional features from real high-dimensional laser welding data. As a second contribution and novel integration step, these features are then used as input to a temporal-difference learning algorithm (in this case a general-value-function learner) to acquire important real-time information about the process of laser welding; temporally extended predictions are used in combination with deep learning to directly map sensor data to the final quality of a welding seam. As a third contribution and final part of our proposed architecture, we suggest that deep learning features and general-value-function predictions can be beneficially combined with actor–critic reinforcement learning to learn context-appropriate control policies to govern welding power in real time. Preliminary control results are demonstrated using multiple runs with a laser-welding simulator. The proposed intelligent laser-welding architecture combines representation, prediction, and control learning: three of the main hallmarks of an intelligent system. As such, we suggest that an integration approach like the one described in this work has the capacity to improve laser welding performance without ongoing and time-intensive human assistance. Our architecture therefore promises to address several key requirements of modern industry. To our knowledge, this architecture is the first demonstrated combination of deep learning and general value functions. It also represents the first use of deep learning for laser welding specifically and production engineering in general. We believe that it would be straightforward to adapt our architecture for use in other industrial and production engineering settings.

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 model-based approach 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

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

http://dx.doi.org/10.1016/j.mechatronics.2015.09.004
0957-4158/

important aspects of intelligence during welding, namely perception, prediction, and interaction.

Representation: As a laser-welding system's 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 transformation-invariant representation of the systems 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 auto-encoders have been shown to successfully compete with state-of-the art feature extraction techniques (e.g., principal component analysis, linear discriminant analysis) [7] and improved [8] or directly learned [9,10] policies for high-dimensional image data in reinforcement learning. Stacked denoising auto-encoders have shown the capability of achieving a general representation, which leads to more robustness against varying data and overfitting [11].

Prediction: 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, envelope curves or look-up tables. But these techniques are restricted either in applicability (a priori model), accuracy (envelope curves) or scalability (look-up tables). They are also limited in their capability to adapt to changes. To deal with these issues, we include predictions about process quality and state as an important part of intelligence [12] in our architecture; importantly, we suggest that predictions should be able to be learned and adapted during the ongoing operation of a system. To date, prediction learning has been dominated by linear models that are difficult to apply to nonlinear and time-varying problems [13]. These problems have been overcome by recent research using the temporal-difference (TD) reinforcement learning approach [14]. New techniques have extended classical TD-learning to allow generalized online predictions [15]. We include these predictions into our proposed system using a temporally extended prediction approach called nexting [16] with general value functions, an approach that is capable of learning and making real-time predictions at multiple timescales.

Control: 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 setup 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 [17]. 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 computation can be done incrementally (linearly) and they can be updated within milliseconds. Due to the fact that experience—from which the algorithm already had learned—does not need to be stored, the memory requirements do not increase over time [18]. By using function approximation they also scale well to real world problems; this has been shown in various applications [19–22].

Our proposed architecture [23] 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 auto-encoders 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 [24]. Despite the environmental uncertainties that the process is exposed to, like changes in temperature, humidity or 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. Fig. 1 illustrates examples for laser welds with different quality.

In our setting, process monitoring is done by a camera-based system and photodiodes, which is a common setting in laser welding applications [25]. As the keyhole, which is the area where the laser hits the material, oscillates with a typical frequency of 500 Hz [26], 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 [27] with a resolution of 144×176 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 1100 nm and 1800 nm. The second observes the plasma radiation at a wavelength of 400–600 nm. The third diode records the laser back reflection at 1050–1080 nm.

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 [28].

Fig. 2 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 has the capacity 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.