

AN IN-DEPTH REVIEW OF THE SCIENTIFIC RESEARCH ON DEEP REINFORCEMENT LEARNING IN PRODUCTION SYSTEMS

Abstract

Production systems face significant problems as a result of shorter product development cycles and fully customizable goods. These must not only handle a greater variety of products, but also facilitate high throughputs and offer a high degree of flexibility and resilience to process changes and unanticipated events. Deep Reinforcement Learning (RL) has been used more and more for production system optimization in order to overcome these obstacles. Deep RL, in contrast to other machine learning techniques, uses freshly gathered sensor data in close contact with its surroundings to allow for real-time reactions to system modifications. A thorough review of the outcomes has not yet been established, despite the fact that deep RL is now being implemented in production systems. This paper's primary contribution is to give practitioners and researchers a summary of applications.

1. Introduction

Companies nowadays have to deal with mass customization and shortened development cycles, which present significant obstacles for intelligent manufacturing facilities. In addition to meeting the ever-tougher requirements for product quality and sustainability in the shortest amount of time, they must be able to function in extremely unpredictable weather. In 2013, the German government initiated the Industry 4.0 program to encourage the creation of adaptable and flexible production systems in order to address these issues (Kagermann, Walter, and Helbig Citation 2013). The initiative has enormous potential and potential impact, however according to Xu, Xu, and Li (Citation 2018), many of the current Industry 4.0 implementations do not yet use correspondingly advanced technology like machine learning. Additionally, Liao et al. (Citation 2017) make this clear when they say that whether modeling, virtualization, or big data techniques

Reinforcement learning, unsupervised learning, and (semi-)supervised learning are all included in the field of machine learning. A (pre-labelled) collection of data is necessary for both supervised and unsupervised learning, but RL is unique in that it involves learning by direct contact with its surroundings. It can adapt flexibly to unpredictable situations and learns by trial and error without the need for any previously gathered data or prior (human) knowledge (Sutton and Barto Citation 2017). Our research attempts to represent the current state-of-the-art of real or simulated deep RL implementations in production systems, taking into account these adaptable and desired properties in contemporary production. In addition, we want to pinpoint current issues and contribute to the definition of future study areas.

As far as we are aware, this is the first effort to document broad uses of deep reinforcement learning in real-world systems. In order to help researchers discover potential future applications and research areas in deep reinforcement learning, we plan to present a systematic overview of current work. The review encourages practitioners to apply research findings to real-world systems and helps them think through potential deployment scenarios. We try to respond to the following research questions in order


RQ1: In which production system domains are deep reinforcement learning applications found?

RQ2: What are the present obstacles to deep reinforcement learning application in production systems.

RQ3: What more study is required to tackle the current issues with deep reinforcement learning in production system.

2. Introduction of reinforcement learning

The trial-and-error learning technique in direct interaction with its environment represents the way reinforcement learning (RL), an instance of machine learning, sets itself apart from supervised and unsupervised learning (Sutton and Barto Citation 2017). It is taken into consideration whenever issues need to be treated in dynamic situations that call for a real-time, reaction-driven decision-making process. It does not require supervision or a pre-defined set of data, either labeled or unlabeled. By using sequential decision-making, it allows for online adaptation to changing environmental conditions (as in Palombarini and Martínez (Citation2019)) and can generalize previously gained knowledge (Wang et al. Citation2020). In reinforcement learning, the agent picks up a policy that, based on the received state, produces an action, as shown Figure 1. Agent–environment interaction; Sutton and Barto.



The curse of dimensionality, that results in an exponentially growing table size, high iterative computational costs, low learning efficiencies, and deteriorated performances, is imposed by the necessary action and state discretization in high-dimensional problem spaces. Deep RL aims to overcome this issue by fusing the benefits of deep learning and reinforcement learning, as suggested by Lee and Kim among others. A neural network, which can handle massive volumes of raw and unsorted input data, maps the policy as a function approximator in deep reinforcement learning (Lange, Riedmiller, and Voigtlander Citation).


Deep reinforcement learning, which was once restricted to the Atari platform by Mnih et al. (Citation2013), is now being used in a growing range of applications because of its adaptability and flexibility online. According to Rossit, Tohmé, and Frutos (Citation2019), the collaborative features and distributed multi-agent capabilities have the potential to greatly boost resilience in applications like smart scheduling. Deep reinforcement learning is therefore a promising method to boost the efficiency of contemporary production systems and facilitate the shift to industry 4.0. However, the intersection of deep RL in many production system domains was not explicitly discussed, in contrast to other algorithmic overviews or the generic explanations of machine intelligence applications in production. In order to close this disparity and emphasize the advantages.

3. Methods of research

The fundamental literature review procedure for deep reinforcement learning applications in production systems is described in this section. We adhere to the content analysis standards provided by Tranfield, Denyer, and Smart (Citation 2003) and Antônio Márcio Tavares, Felipe Scavarda, and José Scavarda (Citation 2016) in order to guarantee a methodical and representative assessment. This makes it possible to compile and assess the body of existing literature and gives the current state-of-the-art in the targeted field. In addition to offering research incentives and managerial insights, the consolidation will help researchers and others discover research needs (Petticrew and Roberts Citation 2006).

Antônio Márcio Tavares, Felipe Scavarda, and José Scavarda (Citation 2016) proposed a guideline that states that the systematic literature review (SLR) can be divided into eight (iterative) processes. These key actions are described in order.

Figure 2. Eight step approach to conduct an SLR



3. Database and iterative keyword selection in Phase 1

As Lohmer and Lasch (Citation2020) or other researchers, we used the Web of Science (all areas), ScienceDirect (title, abstract, or author-specified keywords), and IEEE Xplore (journals) as search databases for our review. We had a rather broad emphasis, consisting of an algorithmic, a general, and a more particular topic, and we defined the keywords using an interactive method to guarantee a representative coverage of the study literature. We included assembly, automation, and industry as generic keywords in the iterative process in addition to production and manufacturing. Additional subcategories, such as quality control, maintenance, and others, were included to the search in order to ensure that no sub-discipline was overlooked (see Table 2). Since the phrase "deep RL" isn't always used.

3.1. Phase 2: Establishing criteria for inclusion and exclusion

We established a number of inclusion and exclusion criteria in order to methodically reduce the scope and guarantee a high level of review quality. As stated in Light and Pillemer (Citation1984) and Durach, Kembro, and Wieland (Citation2017), we only took into account articles from peer-reviewed journals, proceedings, conference papers, and books for quality considerations. Pre-prints, working papers, and other non-peer reviewed publications were not included. Additionally, articles written in languages other than English were disqualified, and only papers published after 2010 were included because of the notable achievements of deep reinforcement learning, particularly the publication by Mnih et al. (Citation 2013). The established research questions and taxonomy structure guarantee a conceptual characterization of the inclusion and exclusion criteria. Our goal was to find industrial deep RL applications, hence we eliminated articles that concentrated on

4. Conclusion

It became understood that deep RL is used extensively in a variety of fields, including maintenance and process control, and that it typically outperforms conventional algorithms because to its adaptability to a wide range of scenarios and ability to handle current production uncertainties (RQ1). In addition to lowering lead times and work-in-progress, achieving high assembly accuracy, or creating strong scheduling policies, this also lessened the present disadvantages of traditional approaches, like their limited capacity for adaptation, high reliance on human judgment, or costly re-optimization.

5. References

1. Altenmüller, Thomas, Tillmann Stüker, Bernd Waschneck, Andreas Kuhnle, and Gisela Lanza. 2020. "Reinforcement Learning for An Intelligent and Autonomous Production Control of Complex Job-Shops Under Time Constraints." *Production Engineering* 14 (3): 319–328.
2. Andersen, Rasmus E., Steffen Madsen, Alexander B. K. Barlo, Sebastian B. Johansen, Morten Nør, Rasmus S. Andersen, and Simon Bøgh. 2019. "Self-Learning Processes in Smart Factories: Deep Reinforcement Learning for Process Control of Robot Brine Injection." *Procedia Manufacturing*.
3. Antônio Márcio Tavares, Thomé, Luiz Felipe Scavarda, and Annibal José Scavarda. 2016. "Conducting Systematic Literature Review in Operations Management." *Production Planning & Control* 27 (5):
4. Arinez, Jorge F., Qing Chang, Robert X. Gao, Chengying Xu, and Jianjing Zhang. 2020. "Artificial Intelligence in Advanced Manufacturing: Current Status and Future Outlook." *Journal of Manufacturing Science and Engineering* 142 (11).
5. Baer, Schirin, Jupiter Bakakeu, Richard Meyes, and Tobias Meisen. 2019, September. "Multi-Agent Reinforcement Learning for Job Shop Scheduling in Flexible Manufacturing Systems." 2019 Second International Conference on Artificial Intelligence for Industries (AI4I). Laguna Hills, CA: IEEE, pp. 22–25.
6. Baer, Schirin, Danielle Turner, Punit Mohanty, Vladimir Samsonov, Romuald Bakakeu, and Tobias Meisen. 2020. "Multi Agent Deep Q-Network Approach for Online Job Shop Scheduling in Flexible Manufacturing." 2020 International Conference on Manufacturing System and Multiple Machines, Tokyo, Japan, pp. 1–9.
7. Bakakeu, Jupiter, Schirin Baer, Jochen Bauer, Hans-Henning Klos, Jörn Peschke, Adrian Fehrle, Werner Eberlein, et al. 2018. "An Artificial Intelligence Approach for Online Optimization of Flexible Manufacturing Systems." *Applied Mechanics and Materials* 882: 96–108.
8. Bakakeu, Jupiter, Dominik Kisskalt, Joerg Franke, Shirin Baer, Hans-Henning Klos, and Joern Peschke. 2020, August 30–September 2. "Multi-Agent Reinforcement Learning for the Energy Optimization of Cyber-Physical Production Systems." 2020 IEEE Canadian Conference on Electrical and Computer Engineering (CCECE), London, ON, Canada. IEEE, pp. 2–8.
9. Bellman, Richard. 1957. *Dynamic Programming*. 1 vols. Princeton: Princeton University Press.
10. Beltran-Hernandez, Cristian C., Damien Petit, Ixchel G. Ramirez-Alpizar, and Kensuke Harada. 2020. "Variable Compliance Control for Robotic Peg-in-Hole Assembly: A Deep-Reinforcement-Learning Approach." *Applied Sciences* 10 (19): 6923.
11. Beltran-Hernandez, C. C., D. Petit, I. G. Ramirez-Alpizar, T. Nishi, S. Kikuchi, T. Matsubara, and K. Harada. 2020. "Learning Force Control for Contact-Rich Manipulation Tasks With Rigid Position-Controlled Robots." *IEEE Robotics and Automation Letters* 5 (4): 5709–5716.