



Contents lists available at ScienceDirect

Annual Reviews in Control

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

Vision article

Advancing systems and control research in the era of ML and AI

Pramod P. Khargonekar*, Munther A. Dahleh

^a Department of Electrical Engineering and Computer Science, University of California, Irvine, USA^b Department of Electrical Engineering and Computer Science Institute for Data, Systems, and Society MIT, USA

ARTICLE INFO

Article history:

Received 26 February 2018

Revised 25 March 2018

Accepted 2 April 2018

Available online xxx

Keywords:

Machine learning

Artificial intelligence

Control systems

ABSTRACT

Fields of machine learning and artificial intelligence are undergoing transformative advances and growth. This article presents a vision for the field of systems and control that simultaneously leverages these advances to more fully engage with them and spur new expansive research directions in systems and control.

© 2018 Elsevier Ltd. All rights reserved.

Contents

References 4

How might the current and expected future advances in machine learning and artificial intelligence lead to new opportunities for the systems and control community? We have been motivated by this timely question to articulate an initial vision in this brief essay. Right at the outset, we would like to state that the discussion below is far from comprehensive. Rather, the purpose is to present a perspective and some initial thinking on how the systems and control community could engage in and help shape this emerging future more fully.

Machine learning and artificial intelligence context:

Since its inception, the field of artificial intelligence (AI) focused on understanding how computers can mimic human brains in the context of decision making. Coined by John McCarthy in 1956, AI is a very general field covering disparate topics such as search, planning, reasoning, learning, natural language processing, perception, vision, etc. Machine learning, coined by Arthur Samuel, focuses on achieving AI through training and learning, and has its roots in statistical learning theory. Over the last decade, great advances have been made in AI and ML in several dimensions: theory, application, and implementation.

Specifically in the context of ML, deep learning architectures, algorithms and techniques have created powerful tools to learn representations of large volumes of data in multiple layers of repre-

sentation (LeCun, Bengio, & Hinton, 2015). These tools appear to be very successful in learning complex functions and discovering intricate structures in high-dimensional data. They have shown superior performance in image and speech recognition and are being applied in a wide variety of problems: drug discovery, particle physics, astronomy, and biomedicine. While it is difficult to summarize (and beyond the purpose and scope of this paper,) all advancements in ML, there are several directions that are particularly relevant to the subsequent discussion.

1. High dimensional statistics that focuses on computational and statistical issues pertaining to learning high dimensional sparse sets of parameters from observations (Tropp, 2015). This development ties well with research areas in compressive sensing, but expands the results to include more elaborate sparse models such as sparse graphs.
2. Online learning which addresses sequential learning and decision making in bandit problems (Rakhlin, Sridharan, & Tewari, 2015). Developments in this direction resulted in different algorithms under different information structures that can guarantee asymptotic optimality. Regret-based algorithms are one outcome of these developments.
3. Discovery of structural and latent variables using spectral and tensor methods (Anandkumar, Ge, Hsu, Kakade, & Telgarsky, 2014). Such an approach became popular in the context of topic modeling and mixtures of Gaussian models where the underlying mixture is not a priori known.

* Corresponding author.

E-mail address: pramod.khargonekar@uci.edu (P.P. Khargonekar).

4. Optimization for ML including the exploitation of randomness in the context of maximum likelihood learning for non-convex network models (Zwiernik, Uhler, & Richards, 2017) and stochastic gradient algorithms in the context of deep learning (Hardt, Recht, & Singer, 2016; Rakhlin, Raginsky, & Telgarsky, 2017).

Advancements in these areas have substantially influenced many application domains.

More broadly, the confluence of insights and techniques from neuroscience, cognitive science, reinforcement learning (RL), and deep learning (DL) has led to very impressive progress in ML and AI with amazing achievements in championship games and demonstration of human level control by an artificial agent (Hassabis, Kumaran, Summerfield, & Botvinick, 2017; Mnih et al., 2015). Increasing computational power (thanks to Moore's Law progress), availability of large amounts of data, and development of more effective algorithms have been critical to many of these successes and will be increasingly even more important to continuing progress (Dean, Patterson, & Young, 2018; Simonite, 2017). Driven by the excitement of progress and potential for major benefits, there is a flood of interest, and corresponding investments, in ML and AI from the academic, industrial and government sectors (Bughin et al., 2017). Yet, there are significant weaknesses and issues that need to be resolved for future progress. These include insufficient and incomplete theoretical foundations, need for large amounts of data, lack of robustness and vulnerability to adversarial attacks, lack of transparency and explainability, biases resulting from algorithms and data, etc.

It is our position that AI should entail some sort of learning combined with decision making. However, in many application areas, AI refers to decision making more broadly. For many of these AI systems, their ability to learn, especially in real-time, is quite limited. As Brooks observes in his very interesting critique of exaggerated AI predictions (Brooks, 2017), "Today's machine learning is not at all the sponge-like learning that humans engage in, making rapid progress in a new domain without having to be surgically altered or purpose-built. ... When humans play a game, a small change in rules does not throw them off. Not so for AlphaGo or Deep Blue." While the latest AI systems such as AlphaGo Zero (Silver et al., 2017) exhibit certain impressive learning abilities, it is far from clear whether these systems can learn to adapt to and deal with rapid and unforeseen changes in the environment.

Control systems: Control systems have a deep, broad and strong base of foundational knowledge developed over the last 60 years with major emphasis on decision making under uncertainty. Dynamic systems modeling, structural properties, model reduction, identification, stability, feedback, optimality, robustness, adaptation, fault tolerance, and architecture have been among the central concerns on the theoretical side. These issues have been explored in a wide variety of settings: linear, nonlinear, stochastic, hybrid, distributed, supervisory, and others. Applications have been wide ranging: aerospace, automotive, manufacturing, chemical process, energy, power, transportation, etc. While there is a very rich history, the future is just as promising as there are a multitude of directions for future theoretical and applications research (Lamnabhi-Lagarrigue et al., 2017). Despite all the progress in various subfields of systems and control, much remains to be done to satisfactorily address control of large, complex, distributed dynamical systems under rapid changes in the environment and high levels of uncertainty.

Future This is truly an opportune moment to develop a forward looking vision that can inspire talented researchers for the next decade or longer. On the one hand, a major goal of AI is to build machines that can learn and think for themselves (Lake, Ullman, Tenenbaum, & Gershman, 2017), including having imagina-

tion, reasoning, planning, etc. On the other hand, we have a rich body of knowledge in control systems. The field of control can both benefit from and influence the ongoing revolutionary advances in ML and AI. These advances in ML and AI are going to be driven by the large increases in computation power and data, intense interest across the world, and large investments in these fields across academic, industrial and government organizations. By leveraging these ongoing trends and advances in ML/AI, we can aim to have significantly more powerful and versatile control systems. For this, we would need to define specific goals that are currently unachievable with existing control techniques but could potentially be achieved by leveraging ML/AI advances. Such goals would likely be driven by major application areas for control. They would have implications and opportunities for theoretical developments in control. On the other side, we can identify ideas, tools and techniques from control systems that have the potential to advance AI in its quest of building machines that learn and think for themselves. Examples include principles and techniques from robust and adaptive control, stochastic control, dynamic programming, system identification, model predictive control, decentralized and distributed control, and agent based control.

Of course, there are historic connections between learning, artificial intelligence, and control systems going back to the 60's. Research fields such as intelligent control and neural networks for control arose from these long standing connections. Closer to the recent developments in ML and AI that are the main focus of this paper, there are deep connections between RL (Sutton & Barto, 1998) and stochastic control (Bertsekas & Tsitsiklis, 1995). There are more recent and much less investigated connections between learning in sensorimotor neural systems (Grush, 2004; Wolpert, Ghahramani, & Flanagan, 2001) and controls, e.g., the role of forward and inverse models for control in the central nervous system, emulation theory of representation that builds on control and Kalman filtering, etc. More generally, very recent developments in neuroscience such as the free energy principle and a unified brain theory (Friston, 2010) are deeply connected with central ideas in systems and control but have received not much attention in controls community and where further explorations are likely to be very fruitful.

In a more speculative and somewhat controversial longer-term direction, while there is acceptance within the AI community that rich internal models are critical to human like learning and decision making and that the learning processes must be informed and constrained by prior knowledge, there is considerable debate within the AI community on whether such internal models should be configured by human designers or should be learned by the AI agent *de novo* (Botvinick et al., 2017; Lake et al., 2017). (Such debates are not far from the Chomsky-Skinner debate which articulated a fundamental dichotomy in understanding language acquisition; one side that is based on the learning of a fixed architecture in the brain (Chomsky) and the other that is based on statistical interpretation of the relationships between the past and present behavior (Skinner) (Chomsky, 1959; Skinner, 2014).) Debates around these questions in AI are far from settled and the research field is rich with open questions and may turn out to offer new opportunities, in view of the centrality of model building aspects in the debates, for systems and control community to engage, contribute, and benefit.

It is worthwhile emphasizing that a substantial part of control theory focuses on fundamental limits of stability and performance. For example, the theory addresses questions such as characterizing the minimal information needed about a process to control it (say to achieve a desired objective). At the same time, statistical learning theory (both standard and high dimensional) focuses on obtaining information-theoretic limits of achievable model accuracy from data. Can these two seemingly disparate paradigms be

Download English Version:

<https://daneshyari.com/en/article/7107652>

Download Persian Version:

<https://daneshyari.com/article/7107652>

[Daneshyari.com](https://daneshyari.com)