Chapter 1
Engineering AI Systems:
A Research Agenda

Jan Bosch
https://orcid.org/0000-0003-2854-722X
Chalmers University of Technology, Sweden

Helena Holmström Olsson
Malmö University, Sweden

Ivica Crnkovic
Chalmers University of Technology, Sweden

ABSTRACT

Artificial intelligence (AI) and machine learning (ML) are increasingly broadly adopted in industry. However, based on well over a dozen case studies, we have learned that deploying industry-strength, production quality ML models in systems proves to be challenging. Companies experience challenges related to data quality, design methods and processes, performance of models as well as deployment and compliance. We learned that a new, structured engineering approach is required to construct and evolve systems that contain ML/DL components. In this chapter, the authors provide a conceptualization of the typical evolution patterns that companies experience when employing ML as well as an overview of the key problems experienced by the companies that they have studied. The main contribution of the chapter is a research agenda for AI engineering that provides an overview of the key engineering challenges surrounding ML solutions and an overview of open items that need to be addressed by the research community at large.

INTRODUCTION

The prominence of artificial intelligence (AI) and specifically machine- and deep-learning (ML/DL) solutions has grown exponentially, see Amershi et al. (2019), and Bernardi et al. (2019). Because of the Big Data era, more data is available than ever before, and this data can be used for training ML/DL

DOI: 10.4018/978-1-7998-5101-1.ch001

solutions. In parallel, progress in high-performance parallel hardware such as GPUs and FPGAs allows for training solutions of scales unfathomable even a decade ago. These two concurrent technology developments are at the heart of the rapid adoption of ML/DL solutions.

Virtually every company has an AI initiative ongoing and the number of experiments and prototypes in the industry is phenomenal. Although earlier the province of large Software-as-a-Service (SaaS) companies, our research shows democratization of AI and broad adoption across the entire industry, ranging from startups to large cyber-physical systems companies. ML solutions are deployed in telecommunications, healthcare, automotive, internet-of-things (IoT) as well as numerous other industries and we expect exponential growth in the number of deployments across society. As examples, ML solutions are used in the automotive industry to explore autonomous driving and as a means to increase efficiency and productivity. In domains such as e.g. mining, autonomous vehicles are currently being used in under-ground operations where human safety is a concern and in situations where there is a risk of accidents. Similarly, self-driving trucks can operate largely automatically within e.g. harbor or airport areas which helps to increase both productivity and safety. In the defense domain, AI segmentation is used to identify buildings, roads, or any type of land at pixel level from a great height. Besides, AI technologies provide a range of opportunities in a fast-moving emergency where there is conflicting information and where there is a need to rapidly establish an understanding of the current situation, as well as for prediction of future events.

Across industries, image recognition capabilities are key and as an example from the packaging domain, ML is used for checking the inner sides of packages to detect any flaws or deviations in sealings and for analyzing temperature, anomalies, and edges to ensure quality.

Unfortunately, our research, see Arpteg et al. (2018), Lwaktare et al. (2019), and Munappy et al. (2019), shows that the transition from prototype to the production-quality deployment of ML models proves to be challenging for many companies. Though not recognized by many, the engineering challenges surrounding ML prove to be significant. In our research, we have studied well over a dozen cases and identified the problems that these companies experience as they adopt ML. These problems are concerned with a range of topics including data quality, design methods, and processes, the performance of models as well as deployment and compliance.

To the best of our knowledge, no research exists that provide a systematic overview of the research challenges associated with the emerging field of AI engineering, which we define as follows:

AI Engineering is an engineering discipline that is concerned with all aspects of the development and evolution of AI systems, i.e. systems that include AI components. AI engineering is primarily an extension of Software Engineering, but it also includes methods and technologies from data science and AI in general.

In this chapter, we provide a research agenda that has been derived from the research that we have conducted to date. The goal of this research agenda is to inspire for the software engineering research community to start addressing the AI engineering challenges.

The purpose and contribution of this chapter are threefold. First, we provide a conceptualization of the typical evolution patterns concerned with the adoption of AI that companies experience. Second, we provide an overview of the engineering challenges surrounding ML solutions. Third, we provide a research agenda and overview of open items that need to be addressed by the research community at large.