



Rethinking Computer Science Through AI

Kristian Kersting¹

© The Author(s) 2020

Most of artificial intelligence (AI) in use today falls under the categories of the first two waves of AI research. First wave AI systems follow clear rules, written by programmer, aiming to cover every eventuality. Second wave AI systems are the kind that use statistical learning to arrive at an answer for a certain type of problem. Think of image classification system. The third wave of AI envisions a future in which AI systems are more than just tools that execute human programmed rules or generalize from human-curated data sets. The systems will function as partners rather than as tools. They can acquire human-like communication and reasoning capabilities, with the ability to recognize new situations and to adapt to them. For example, a third wave AI system might note that a speed limit of 120 km/h does not make sense when entering a small village by car.

In my opinion, it is time to usher in the third way of AI. Current second wave AI systems are highly specialized systems that are typically very good at specific, well-defined tasks. They are often not robust and without extensive re-training often fail even for modestly different circumstances. For instance, an object in a non-canonical orientation and context fools many second wave AIs for visual scene understanding. As Gary Marcus points out, they may fail to recognize a school bus tipped over on its side in the context of a snowy road. Understanding such limitations of second wave approaches is, to name only one of many instances, particularly important for self-driving cars. As impressive as they are already at the current stage, they utilize a composite of independent and narrow intelligent subsystems. Following Scott Jones, if you took the software from a self-driving car and put it in a golf cart, it is likely to be useless without considerable re-programming and -training. In contrast, any human who has learned to drive a car could get into a

golf cart for the first time and would have no major problem navigating the fairways. This is because humans are very good at abstraction: we can easily generalize solutions and apply them to similar but different problems. And even if we were encountering problems driving the golf cart, we could articulate them and ask for help.

The third wave of AI poses many deep and fascinating scientific problems: How do we bring together different—and currently separated—AI regimes: low-level perception and high-level reasoning? Akin to Systems Biology, how should a systemic view on AI look like allows us to capture, understands and utilize individual AI algorithms as building blocks for a complex AI system in a mathematical and computationally sound way? How do we manage that non AI experts build, use and interact with AI systems? How do we make machine learning a co-adaptive process, in which a human is changing computer behavior, but the human also adapts to use machine learning more effectively and adapts his or her data and goals in response to what is learned by the machine? How do we increase the efficacy of the “teachers” given the learners, i.e., how to we move from machine learning to machine teaching? How to balance tasks that we would find useful to automate and tasks in which it might remain meaningful for us humans to participate. How do we incorporate computational models for human intelligence into AI systems so that machines can learn so much about the world, so rapidly and flexibly, as humans?

Meeting these challenges not only pushes AI. It provides a unique opportunity to rethink fundamental problems and methods of computer science itself, from hardware and software design, over databases and robotics, to human–computer interaction and software engineering. The third wave of AI is a team sport that holds the chance to grow even closer together in the CS departments and reaching out to other fields such as (computational) cognitive science.

The current special issue on Ontologies and Data Management, its second part, illustrates this very nicely. The contributions underscore the great value of combining reasoning and learning. Enjoy learning about the latest developments!

Stay healthy,

✉ Kristian Kersting
kersting@cs.tu-darmstadt.de

¹ Artificial Intelligence and Machine Learning Group,
Computer Science Department and Center for Cognitive
Science, TU Darmstadt, Hochschulstrasse 1, Room 074,
64289 Darmstadt, Germany

Kristian Kersting.

1 Forthcoming Special Issues

1.1 Developmental Robotics

Guest Editors: Manfred Eppe, Verena V. Hafner, Yukie Nagai, Stefan Wermter

Human intelligence develops through experience, robot intelligence is engineered—is it? At least in the mainstream approaches based on classical artificial intelligence (AI) and machine learning (ML) the robotic engineering approach is pursued and data- or knowledge-based algorithms are designed to improve a robot's problem-solving performance. Based on this engineering perspective of classical AI/ML approaches plenty of valuable application-specific impact has been achieved. Yet, the achievements are often subject to restrictions that involve domain knowledge as well as constraints concerning application domains and computational hardware.

Developmental robotics seeks to extend this constrained perspective of engineered artificial robotic cognition, by building on inspiration from biological developmental processes to design robots that learn in an open-ended continuous fashion. Developmental robotics considers cognitive domains that involve problem-solving, self-perception, developmental disorders and embodied cognition.

This perspective helps to improve the performance of intelligent robotic agents, and it has already led to significant contributions that inspired cutting-edge application-oriented machine learning technology. In addition, developmental robotics also provides functional computational models that help to understand and to investigate embodied cognitive processes.

For this special issue, we welcome contributions that include, but are not limited to the following topics:

Robotic self-perception and body representation; typical development and developmental disorders; neural foundations of development and learning; continual learning; transfer learning; embodied cognition; problem-solving; predictive models; intrinsic motivation; language LEARNING.

2 Education in Artificial Intelligence K-12

Guest Editors: Gerald Steinbauer, Martin Kandlhofer, Tara Chklovski, Fredrik Heintz, Sven Koenig

The upcoming special issue of the KI Magazin addresses the emerging topic of education in artificial intelligence (AI) at the K-12 level. In recent years, artificial intelligence (AI) has attracted a lot of attention from the public, and become a major topic of economic and

societal discussion. AI already has a significant influence on various areas of life and across different sectors and fields. The speed and force with which AI is impacting our work and everyday life poses a tremendous challenge for our society and educational system. Teaching fundamental AI concepts and techniques has traditionally been done at the university level. However, in recent years several initiatives and projects pursuing the mission of K-12 AI education have emerged. In this context we also see education organizations and AI experts as well as governments developing and deploying AI-curricula and programs for a K-12 audience. The aim of this special issue is to provide a compact overview over this growing field. We invite contributions from researchers, practitioners, and educators interested in education in AI at K-12 level.

3 Special Issue: NLP and Semantics


Guest Editors: Daniel Hershcovich, Lucia Donatelli and Stephan Oepen

Making computers as intelligent as humans has been argued to be as difficult as making them understand human language, which is one of the focus points of natural language processing (NLP). The field has been changing over the past decades, generally moving from rule-based methods to statistical ones. Machine learning (ML) methods, in particular deep learning, are today omnipresent, challenging methods based on linguistic theories by fully end-to-end data-driven modeling. However, combining powerful ML models with flexible pipelines and frameworks based on human and linguistic insight is an exciting development promising the best of both worlds.

NLP applications are abundant, and are already changing people's lives, enabling effortless translation, learning and interaction with human-centric systems in robotics and virtual assistants. While many classical NLP problems deal with modeling the surface form of linguistic utterances, general natural language understanding and generation depend on explicit or implicit modeling of semantics, including meaning, communicative intent, and the complex mapping to the linguistic form. Computational semantics is the study of how to automate the process of constructing and reasoning with meaning representations of natural language expressions, which can take many forms, such as continuous vectors or discrete graphs.

For this special issue, we welcome contributions including, but not limited to the following topics: lexical semantics, compositional semantics, cross-lingual semantics, semantic parsing, syntax-semantics interface, semantic role labeling, textual inference, formal semantics, coreference, discourse, reading comprehension, knowledge acquisition, common sense reasoning, summarization,

multimodal semantics, semantic annotation, ethical aspects in semantic representations, under specification, ontologies, sentiment analysis, stylistic analysis, argument mining, and human–robot interaction.



Funding Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.