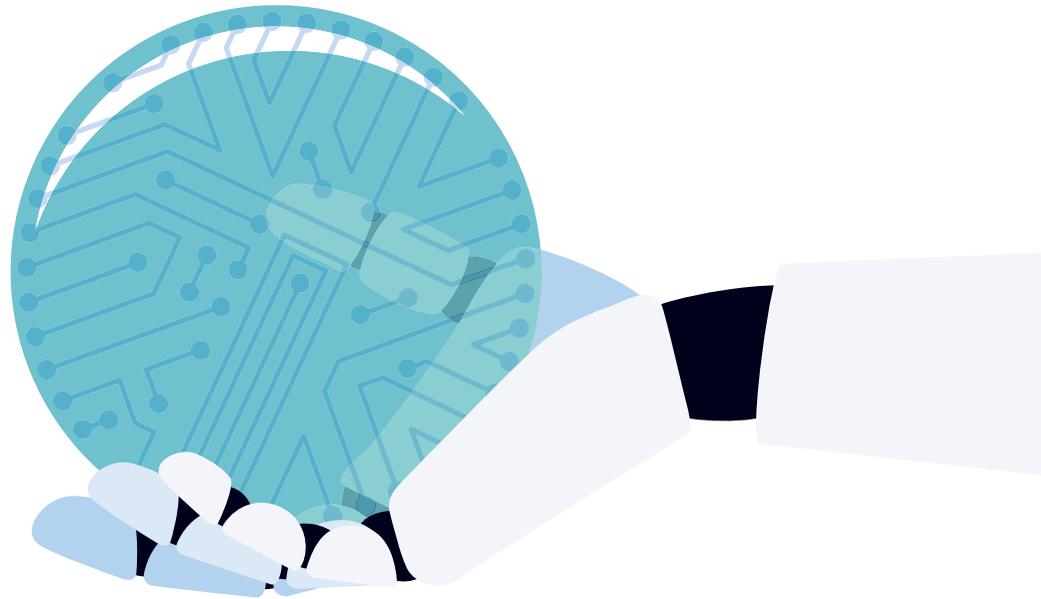


The Expanding Frontier of Artificial Intelligence

Sumi Helal, Lancaster University and
Computer Editor in Chief



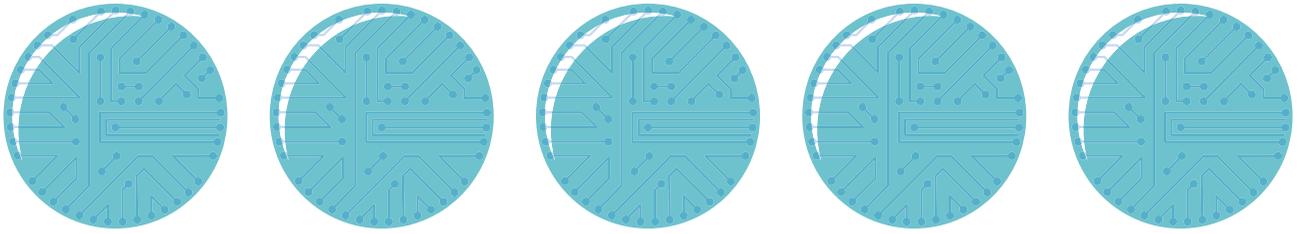


There are many frontiers in computing, each with its own unique profundity in the types of changes it brings forward. With AI, the changes we see will remind us that the playing field for humans and computers is not equal, and how this technology contributes to our lives will present both incredible opportunities as well as some troubling challenges. In this special issue, Computer's editor in chief introduces some of the emerging transformations the future of AI brings to us as humans.

The term “artificial intelligence” was coined a little over six decades ago, but the idea that we could somehow imbue our beloved machines with elements of humanity or the human mind to help us has been with us for much longer. And while there have been many dark imaginings of how AI might pose certain threats to us—think *2001: A Space Odyssey* or even to a far lesser extent Watson’s triumph over human contestants at *Jeopardy*—in 2018, we seem to be much closer to accepting and embracing AI technologies as they evidently demonstrate how they enhance our smartphones, online shopping services, ability to search and find what we like, our smart watches, our health, how we travel, and how businesses connect with and understand customers. The list goes on and on.

Many challenges remain to be met, however, for AI to take off with all on board. First, of course, the massive data required for training AI algorithms is a harsh requirement that is expensive and time-consuming. Also, the more the data we are required to generate and handle for training purposes, the higher the risk for security and privacy breaches. However, perhaps the biggest challenge is in understanding how AI works and how it makes decisions—for the lay user, the expert, the lawyers, the policy makers, the media’s hype cycle, among others. Explaining AI decisions may one day inform the local law enforcement official and the auto insurance company as to which of two autonomous vehicles is at fault, should they be involved in a rare accident. After all, not all autonomous vehicles will be equal.

GUEST EDITOR'S INTRODUCTION



Researchers are taking these challenges head on—working to improve our ability to understand how AI works, our trust and confidence in it, and looking for opportunities for AI to handle the kinds of problems that humans do not enjoy solving.

To give you a sampling of where AI research is taking us, I am pleased to introduce this special issue in which *Computer* explores the future opportunities and challenges of AI from the several experts in this field.

IN THIS ISSUE

In “Toward Anthropomorphic Machine Learning,” Plamen P. Angelov and Xiaowei Gu introduce the ambitious concept of anthropomorphic machine learning as an emerging direction for the next generation of AI. Anthropomorphic machine learning is the idea that machines may be able to learn as humans do. The authors cover the current limitations of mainstream machine-learning techniques (such as deep learning neural networks), including ad hoc embedded structures, a rigid and data-intensive training process that cannot be applied in real-time and evolving systems, and an inability to cope well with the various uncertainties. The authors discuss these limitations in the context of indiscriminate media hype that cites only the promise of AI, while ignoring the serious challenges it faces. The proposed anthropomorphic machine learning approach, which is based on the recently introduced deep rule-based systems, is able to learn continuously by building models incrementally starting from very little training data, learning further as the system evolves. This is a huge departure from the machine learning methods practiced today. Further, this learning model utilizes “prototypes”

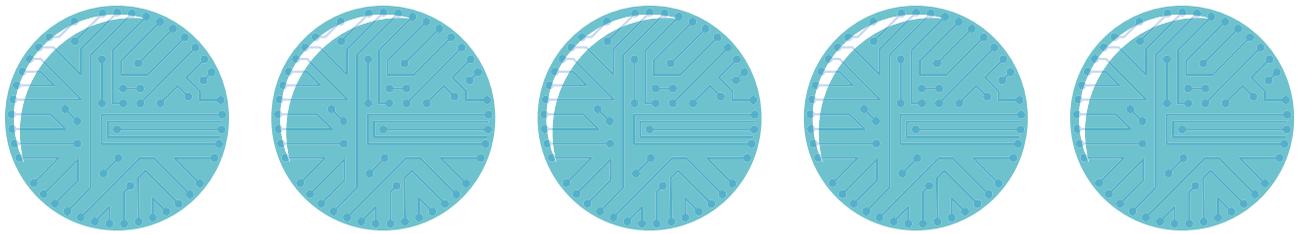
to describe or annotate the learned observations. This is also another huge difference, as prototypes could explain why certain decisions were made and the rationale for learning one thing but not another—just as humans learn. The authors demonstrate this new approach using a deep rule-based image classifier.

In “Toward Human-Understandable, Explainable AI,” Hani Hagraas focuses on the issue of transparency in the decision making of mainstream machine learning algorithms (for example, deep learning), which is currently lacking. Unexplainable and non-auditable “black-box” learning raises concerns from governments and consumers, which hampers meaningful widespread adoption. The author presents three ongoing approaches to enabling explainable AI (XAI), which emerged from the XAI project from the US Defense Advanced Research Project Agency (DARPA). The author then presents a fourth approach based on fuzzy logic systems that attempts to mimic human thinking by being on a continuum of approximations rather than precise or quantized thinking. Specifically, recent revolutionary advances in fuzzy rule-based systems (FRBSs) were able to generate shorter and fewer *if-then* rules (independent of the number of inputs) that retain their clarity and interpretability, which is in contrast to traditional FRBSs. Similar to the first article by Angelov and Gu, Hagraas emphasizes that XAI is a research problem of highest priority and a key challenge that must be resolved satisfactorily for the AI to gain wide acceptability.

In “The Age of Artificial Emotional Intelligence,” Dagmar Schuller and Björn Schuller address the application of AI in the area of emotional

intelligence—which is defined as the ability to recognize emotions, generate and adapt to emotions and affects, and apply emotional informatics in optimizing goal accomplishment or problem solving. The authors show how early artificial emotional intelligence (AEI) employed machine learning algorithms, including hidden Markov models, neural networks, among others, to recognize human emotions in terms of changes in facial expressions, voice acoustics, spoken words, and other modalities. Further, the authors argue that deep learning is currently positioned to overcome the huge variability in recognition performance that results from traditional machine learning techniques. The authors also demonstrate the smaller role AI plays in generating emotions, which amounts to a rule-based system, and the larger role AI plays in emotion augmentation and application into planning, reasoning, and goal achievements. In addition, they describe another augmentation in which emotion-inspired principles are embedded in the machine-learning algorithms—a fascinating concept. An example is given by the authors to help explain emotion augmentation in reinforcement learning, as well as in cognition and abstraction learning. The article lays down the broad vision of how AEI can enable more natural and efficient human-computer interactions, and points to missing pieces of the puzzles and the work ahead to enable this vision.

In “AI and Blockchain: A Disruptive Integration,” Thang N. Dinh and My T. Thai shine the spotlight on an accidental symbiosis between AI and the emerging blockchain technology. As the authors point out, on the one hand, blockchain is a technology that can support AI and its future



advancements. The authors point to how AI is essentially enabled by big data and how such data is difficult to obtain in some application areas (for example, personal health data) due to privacy concerns. Expanding users' rights to privacy to also include copyright (right of ownership of the data) could encourage data sharing in return for royalties. Such schemes would require blockchain to track the use of individuals' data to enable use-based payment. The authors also point to the recent trend in which AI is democratized by allowing consumers (not the big cloud and cloud analytics companies) to collectively become a blockchain (a distributed analytic platform) by utilizing their unused devices (such as mobile phones, gaming devices, set-top boxes, and other emerging IoT devices) to run AI algorithms and perform analysis. Also, the authors point to XAI and discuss how blockchain can help by adding accountability and

traceability of all decisions made by AI algorithms that could be implemented as and over a blockchain. On the other side of this symbiosis, the authors show how AI supports an increased utility of blockchain technology, which can give users control over their data—including, for instance, what they like or dislike in their social network, if such social networks are implemented using blockchains. However, personalization of content, in this example, would be lost, which is a hefty cost for such anonymization.

AI utilized directly by users over such blockchain systems would then help bring back personalization in a safer way. The article predicts significant innovations and revolutionary societal benefits to come out of this integration between AI and blockchains.

This has been a fun and inspiring special issue, and I thank the authors for their great contributions and hope that you find these articles enjoyable and informative. **E**

ABOUT THE AUTHOR

SUMI HELAL is professor and chair in digital health at Lancaster University. He is a Fellow of IEEE and a member of the IEEE Computer Society Board of Governors. Contact him at s.helal@lancaster.ac.uk.

Call for Articles

IEEE Pervasive Computing

seeks accessible, useful papers on the latest peer-reviewed developments in pervasive, mobile, and ubiquitous computing. Topics include hardware technology, software infrastructure, real-world sensing and interaction, human-computer interaction, and systems considerations, including deployment, scalability, security, and privacy.

Author guidelines:

www.computer.org/mc/pervasive/author.htm

Further details:

pervasive@computer.org
www.computer.org/pervasive

IEEE
pervasive
COMPUTING
MOBILE AND UBIQUITOUS SYSTEMS