FEBRUARY 2020 www.computer.org

IEEE IEEE COMPUTER SOCIETY

Keep Your Career Options Open

Upload Your Resume Today!

Whether your enjoy your current position or you are ready for change, the **IEEE Computer Society Jobs Board** is a valuable resource tool.

Take advantage of these special resources for job seekers:

- JOB ALERTS
- TEMPLATES
- CAREER ADVICE
- RESUMES VIEWED BY TOP EMPLOYERS
- WEBINARS

No matter your career level, the IEEE Computer Society Jobs Board keeps you connected to workplace trends and exciting new career prospects.

www.computer.org/jobs

IEEE COMPUTER SOCIETY

COMPUTING edge

IEEE COMPUTER SOCIETY computer.org • +1 714 821 8380

Printed with inks containing soy and/or vegetable oils

IEEE

IEEE COMPUTER SOCIETY

STAFF

Circulation: ComputingEdge (ISSN 2469-7087) is published monthly by the IEEE Computer Society. IEEE Headquarters, Three Park Avenue, 17th Floor, New York, NY 10016-5997; IEEE Computer Society Publications Office, 10662 Los Vaqueros Circle, Los Alamitos, CA 90720; voice +1 714 821 8380; fax +1 714 821 4010; IEEE Computer Society Headquarters, 2001 L Street NW, Suite 700, Washington, DC 20036.

Postmaster: Send address changes to ComputingEdge-IEEE Membership Processing Dept., 445 Hoes Lane, Piscataway, NJ 08855. Periodicals Postage Paid at New York, New York, and at additional mailing offices. Printed in USA.

IEEE Computer Society Magazine Editors in Chief

Computer
Jeff Voas, *NIST*

Computing in Science & Engineering
Lorena A. Barba (Interim), *George Washington University*

IEEE Annals of the History of Computing
Gerardo Con Diaz, *University of California, Davis*

IEEE Computer Graphics and Applications
Torsten Möller, *Universität Wien*

IEEE Intelligent Systems
V.S. Subrahmanian, *Dartmouth College*

IEEE Internet Computing
George Pallis, *University of Cyprus*

IEEE Micro
Lizy Kurian John, *University of Texas at Austin*

IEEE MultiMedia
Shu-Ching Chen, *Florida International University*

IEEE Pervasive Computing
Marc Langheinrich, *Università della Svizzera italiana*

IEEE Security & Privacy
David Nicol, *University of Illinois at Urbana-Champaign*

IEEE Software
Ipek Ozkaya, *Software Engineering Institute*

IT Professional
Irena Bojanova, *NIST*

FEBRUARY 2020 • VOLUME 6, NUMBER 2

COMPUTING edge

51

Smart Edge: The Effects of Shifting the Center of Data Gravity Out of the Cloud

Software

Artificial Intelligence

Automation

Edge Computing

Departments

Magazine Roundup

The IEEE Computer Society's lineup of 12 peer-reviewed technical magazines covers cutting-edge topics ranging from software design and computer graphics to Internet computing and security, from scientific applications and machine intelligence to visualization and microchip design. Here are highlights from recent issues.

Computer

The Arpanet and Its Impact on the State of Networking

In this article from the October 2019 issue of *Computer*, experts discuss technical lessons from the Arpanet, one of the early computer networks and a direct progenitor of the Internet. The Arpanet was planned and put into initial operation between 1967 and 1971 by the Information Processing Techniques Office (IPTO) of the Advanced Research Projects Agency (ARPA, now known as the Defense Advanced Research Projects Agency (DARPA)).

Computing in Science & Engineering

Track Occupation Detection Based on a Maximum Posterior Probability Model using Multisensor Data Fusion

To address the railway track section occupation detection failure, this article from the November/December 2019 issue of *Computing in Science & Engineering* proposes a maximum posterior probability model that uses multisensor information fusion to detect track occupation. Based on the installation method of the sensor, this model obtains stable base data of occupied track sections and extracts its features, including the train's running velocity, acceleration, direction, occupied area, wheelset axle counting, and track vibration. The maximum posterior probability and logarithm model are then derived by computing the prior probability, the posterior probability, and the conditional joint probability for the features. The judgment of the track occupation is more accurate compared with experience value. The experiments demonstrate that the track occupation detection method can effectively judge the occupation of the train, people, and tool cart. Based on the maximum posterior probability, the Bayes optimal data fusion ratio for a measured parameter in this article reaches 99.9983 percent.

Published by the IEEE Computer Society

IEEE Annals of the History of Computing

The Killer App that Saved the Macintosh

In 1985, Apple was in crisis, and Macintosh and LaserWriter sales were plummeting. Marketing manager John Scull was tasked with saving the LaserWriter. Despite having no staff and facing resistance, Scull marshaled allies across Apple's sales and marketing organizations, its dealer channel, and third-party partners like Aldus and Adobe in order to launch a successful desktop publishing marketing program. By 1988, desktop publishing at Apple had become a billion-dollar business. Read more in the July–September 2019 issue of *IEEE Annals of the History of Computing*.

IEEE Computer Graphics and Applications

BEAMES: Interactive Multimodel Steering, Selection, and Inspection for Regression Tasks

Interactive model steering helps people incrementally build machine-learning models that are tailored to their domain and task. Existing visual analytic tools allow people to steer a single model (such as assignment attribute weights used by a dimension reduction model). However, the choice of model is critical in such situations. What if the model chosen is suboptimal for the task, dataset, or question being asked? What if, instead of parameterizing and steering this model, a different model provides a better fit? This article from the September/October 2019 issue of *IEEE Computer Graphics and Applications* presents a technique to allow users to inspect and steer multiple machine-learning models. The technique steers and samples models from a broader set of learning algorithms and model types. The authors incorporate this technique into a visual analytic prototype, BEAMES, that allows users to perform regression tasks via multimodel steering. This article demonstrates the effectiveness of BEAMES via a use case, and discusses broader implications for multimodel steering.

IEEE Intelligent Systems

Learning from Personal Longitudinal Dialog Data

The authors of this article from the July/August 2019 issue of *IEEE Intelligent Systems* explore the use of longitudinal dialog data for two dialog prediction tasks: next message prediction and response time prediction. They show that a neural model using personal data that leverages a combination of message content, style matching, time features, and speaker attributes leads to the best results for both tasks, with error rate reductions of up to 15 percent compared to a classifier that relies exclusively on message content and to a classifier that does not use personal data.

IEEE Internet Computing

Edge-Based Live Video Analytics for Drones

Real-time video analytics on small autonomous drones poses several difficult challenges at the intersection of wireless bandwidth, processing capacity, energy consumption, result accuracy, and timeliness of results. In response to these challenges, this article from the July/August 2019 issue of *IEEE Internet Computing* describes four strategies to build adaptive computer-vision pipelines for domains such as search-and-rescue, surveillance, and wildlife conservation. The experimental results show that a judicious combination of drone-based processing and edge-based processing can save substantial wireless bandwidth and thus improve scalability, without compromising result accuracy or latency.

IEEE Micro

A Hardware–Software Blueprint for Flexible Deep-Learning Specialization

This article from the September/October 2019 issue of *IEEE Micro* describes the Versatile Tensor Accelerator (VTA), a programmable deep-learning architecture designed to be extensible in the face of evolving workloads. VTA achieves "flexible specialization" via a parameterizable architecture, two-level Instruction Set Architecture (ISA), and a Just-in-Time (JIT) compiler.

IEEE MultiMedia

Hierarchical Deep Cosegmentation of Primary Objects in Aerial Videos

Primary object segmentation plays an important role in understanding videos generated by unmanned aerial vehicles. In this article from the July–September 2019 issue of

IEEE MultiMedia, the authors propose a large-scale dataset with 500 aerial videos and manually annotated primary objects. From this dataset, they find that most aerial videos contain large-scale scenes, small primary objects, and consistently varying scales and viewpoints. The authors propose a hierarchical deep cosegmentation approach that repeatedly divides a video into two sub-videos formed by the odd and even frames, respectively. In this manner, the primary objects shared by sub-videos can be cosegmented by training two-stream CNNs and finally refined within the neighborhood reversible flows. Experimental results show that the approach remarkably outperforms 17 state-of-the-art methods in segmenting primary objects in various types of aerial videos.

IEEE Pervasive Computing

Revolution or Evolution? Speech Interaction and HCI Design Guidelines

The evolution of designing interactive interfaces has been rather incremental over the past few decades, largely focused on graphical user interfaces (GUIs), even as these extended from the desktop to mobile or to wearables. Only recently can we engage in ubiquitous, ambient, and seamless interactions, as afforded by voice user interfaces (VUIs) such as smart speakers. The authors of this article from the April–June 2019 issue of *IEEE Pervasive Computing* posit that recent speech engineering advances present an opportunity to revolutionize the design of voice interactions. Yet current design guidelines or heuristics are heavily oriented towards GUI interaction, and thus may not fully facilitate the design of VUIs. The authors survey current research revealing the challenges of applying GUI design principles to this space, as well as critique efforts to develop VUI-specific heuristics. They use these to argue that the path toward revolutionary new ubiquitous conversational voice interactions must be based on several evolutionary steps that build VUI heuristics off existing GUI design principles.

IEEE Security & Privacy

Cyber-physical Security for the Masses: A Survey of the Internet Protocol Suite for Internet of Things Security

Internet of Things (IoT) deployments expand as IoT security lags. This article from the September/October 2019 issue of *IEEE Security & Privacy* surveys IoT security protocols standardized by the Internet Engineering Task Force and discusses remaining gaps. Although these standardized IoT security protocols do not completely secure IoT devices, they go a long way.

IEEE Software

Who Can Maintain This Code?: Assessing the Effectiveness of Repository-Mining Techniques for Identifying Software Maintainers

In large and complex systems, identifying developers capable of maintaining a piece of code is an essential task. Repository-mining techniques can help by providing some level of automation; however, whether such techniques effectively identify skilled software maintainers is still unclear. Read more in the November/December 2019 issue of *IEEE Software*.

IT Professional

Purchase-Based Analytics and Big Data for Actionable Insights

The trend of mining customer loyalty data for insights on consumer purchasing behavior has been in the making of more than three decades of market testing. Various tactics, such as card-linking, have enabled merchants and advertisers to close the gap between digital ads and local in-store paper coupons. Card-linking allows consumers to link or attach their existing credit or debit cards to rewards systems, such as loyalty programs, digital coupons, or non-reward-based systems, such as mobile wallets. While card-linking has been growing within the larger context of the e-commerce and retail sales ecosystem, e-commerce accounted for only 9.8 percent of all US retail sales as of the third quarter of 2018, according to the US Census Bureau (2018). Read more in the September/October 2019 issue of *IT Professional*.

Values in AI Software

Artificial intelligence (AI) attempts to imitate human behavior and reasoning, but what about ethics and morality? As AI becomes more integrated into our lives, there is a growing consensus about the importance of instilling human values such as fairness and respect in AI software. This issue of *ComputingEdge* focuses on strategies for creating ethical AI software.

In *Computer*'s "From Artificial Intelligence to Artificial Wisdom: What Socrates Teaches Us," the authors propose applying Socratic principles to AI software development in the hopes of promoting healthier democracies. "Think Your Artificial Intelligence Software Is Fair? Think Again," from *IEEE Software*, challenges software engineers to use bias-mitigation workflows and tools to help eliminate unjust AI algorithms and models.

AI has many real-world applications in modern society and has the potential for even more impact as the technology improves. *IEEE Security & Privacy*'s "Artificial Intelligence for Law Enforcement: Challenges and Opportunities" identifies explainability, bias handling, and support as key areas that need improvement in order for AI to benefit law enforcement. "Robot Science Writers," from *Computing in Science & Engineering*, discusses AI's ability to produce simple news articles and scientific papers.

AI and other emerging technologies are driving industry automation. *IT Professional*'s "Empowering Extreme Automation via Zero-Touch Operations and GPU Parallelization" posits that zero-touch provisioning and GPU-based computing could bring scalable performance to smart manufacturing. Another article from *IT Professional*, "Security and Vulnerability of Extreme Automation Systems: The IoMT and IoA Case Studies," examines automation in Internet of Things (IoT) applications.

Another trend in IoT is edge computing. *IEEE Internet Computing*'s "Going Back to the Roots—The Evolution of Edge Computing, an IoT Perspective," provides an overview of edge computing—why it was created, what it looks like, and how it's shaping the future of IoT. Finally, *Computer*'s "Smart Edge: The Effects of Shifting the Center of Data Gravity Out of the Cloud" describes the trend of computational resources moving to the edge and fog layers.

From Artificial Intelligence to Artificial Wisdom: What Socrates Teaches Us

Tae Wan Kim, Carnegie Mellon University
Santiago Mejia, Fordham University

A critical examination of existential questions may lead developers to design machines with higher forms of artificial intelligence. Infused by their ability to recognize their own ignorance, these machines would display not merely intelligence but wisdom.

Engineers, especially engineering students, should have an opportunity to think deeply about the nature of human flourishing and human excellence if they want to be educated to develop *good* artificial intelligence (AI). The conventional approach seeks to design AI that avoids causing harm. But this approach falls short to the extent that it does not engage with the question "What is an excellent, flourishing, human being?" Socrates taught us two important things about this question: 1) reflecting on it was a central part of being human and 2) seriously engaging with this question leads to the recognition of a particular form of ignorance that is also a form of wisdom.

In this article, we will elaborate on these Socratic insights and show how they bear on AI. We hope that current and future engineers will be moved to build upon the ancient wisdom discussed here to reimagine their work on AI.

VALUE ALIGNMENT AND HUMAN FLOURISHING

Businesses increasingly use AI to make important decisions for humans. Amazon, Google, and Facebook choose what users see. The driver-assist technology used in most brand-new vehicles aids drivers with steering and braking. Uber and Lyft match passengers with drivers and set prices. Though each of these examples comprises its own complicated technology, they share a core: a data-trained set of decision rules (often called machine learning or AI) that implements a decision with little or no human intermediation.

Digital Object Identifier 10.1109/MC.2019.2929723
Date of publication: 24 September 2019

 Published by the IEEE Computer Society

EDITORS
HAL BERGHEL University of Nevada, Las Vegas; hlb@computer.org
ROBERT N. CHARETTE ITABHI Corp.; rncharette@ieee.org
JOHN L. KING University of Michigan; jlking@umich.edu

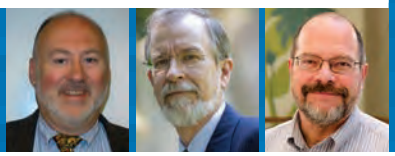

AI techniques are quickly being adopted to automate decisions. As this happens, societal worries about the compatibility of AI and human values grow. How can we ensure that AI does not turn against us? That it is under our control? That it serves us and promotes what we value? In response to these worries, some computer scientists have suggested that "value alignment" should be one of the top priorities in AI research.[1,2] Value alignment seeks to ensure that the technology we design incorporates the values that are important to us. The concept dates back to Alan Turing, who wrote about the need for machines to adapt to human values: "[T]he machine must be allowed to have contact with human beings in order that it may adapt itself to their standards."[3]

The idea of value alignment is consistent with the IEEE's approach to ethics. Recently, the IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems released an ambitious document outlining directives for "ethically aligned design."[4] This document goes beyond the conventional approach that places the no-harm principle as a side constraint on design, emphasizing, instead, that human well-being and human flourishing should be central aims that technology should promote. Some technology developers have lost sight of this too often in recent years, often because their organizations have been too focused on the pursuit of short-term profits and bigger market share. Some of the most important problems to which technology has given rise, and which have turned public opinion distrustful of technological innovation, might have been averted by giving a more prominent role to human well-being and flourishing in the development of such innovations.

Putting human flourishing at the center of value alignment, however, is not simple. Offering a concrete and well-developed account of the nature of flourishing can be seen as a fool's errand. It is always challenging to offer such an account, but it is especially difficult to do so in a world with rapidly evolving technologies. Technologies are designed to solve a variety of human problems. In doing so, however, they inevitably reshape the ways in which humans interact with the world and flourish in it. The invention of the bow and arrow, for instance, enabled humans to hunt from a safe distance. This enabled them to reduce the risks of hunting at close range and expanded the availability of wild game. But the bow and arrow also modified the nature of hunting, thereby redefining what it meant to flourish as a hunter (and, given its application to war, to excel as a warrior). In sum, the difficulty in defining "flourishing" is not merely that it requires clarity about a host of central human concepts that are difficult to pin down but that it is in flux as technology opens and forecloses different kinds of existential possibilities.[5]

Is there any value to reflecting on human flourishing, given these difficulties? Socrates, the ancient figure, helps us to see why the answer is a resounding yes. He teaches us that recognizing that we fall short in articulating the nature of flourishing is a fundamental form of human wisdom. We propose that this form of Socratic wisdom should play a more prominent role in the development of AI.

SOCRATIC IGNORANCE

During the trial at which he was condemned to death, Socrates explained how he had come to be "Athens's gadfly." An impulsive friend of his, Chaerephon, had asked the oracle of Delphi whether there was anyone wiser than Socrates. The oracle replied that no one was wiser. This response puzzled Socrates because he did not consider himself wise; he was aware that he did not have a well-worked-out account of the nature of human flourishing.

In an attempt to clarify the oracle's meaning, Socrates sought those who were reputedly wise and asked them about their wisdom. He talked with politicians, poets, and craftsmen. After examining them through questions aimed to clarify their views and single out their implications, Socrates always reached the same conclusion: "...[N]either of us knows anything worthwhile, but he thinks he knows something when he does not, whereas when I do not know, neither do I think I know. I am likely to be wiser to this small extent, that I do not think I know what I do not know."[6] What Socrates thought was a form of ignorance turned out to be a form of wisdom.

Socratic wisdom, that is to say, Socratic ignorance, brings an increased openness and humility with respect to how the most important human questions should be answered. Socrates wanted to become wiser and did so by conversing with anyone about human flourishing, regardless of age, class, social status, and geographical origin. He did not exclude any view, no matter how apparently outrageous. Instead, he rigorously examined it in the hopes of learning from it. The fact that he was willing to examine everyone and that he was open to all sorts of opinions makes his approach a powerful tonic against echo chambers and filter bubbles. In addition, cultivating Socratic ignorance seems particularly important in a society like ours, where globalization is causing diverse cultures to clash and where technology is redefining, at a very fast pace, what it means to flourish as a human.

SOCRATIC ENGINEERS

AI is a systematic approach to replicating human intelligence by using various mathematical, computational, and mechanical principles. The Turing test originated from "the imitation game," in which a man attempted to replicate a woman's behavior to deceive an

interrogator sitting in a different room.[7] Because AI is meant to imitate human intelligence, it would be worthwhile to reflect on what a Socratic human—a Socratic engineer, to be precise—might look like.

Many engineers suffer from one of two moral ailments. On the one hand, engineers working on narrowly construed technical projects hold the view that technological tools have no moral valence because they are mere instruments. Some engineers believe that because they don't tell people how to use these tools, they are not responsible for how such tools are used. Consequently, it is frequent for those whose work is narrowly defined to think that questions about human flourishing are detachable from their professional tasks, that it is not their place to think about them.

On the other hand, engineers who have successfully developed innovations that have had a significant impact in the world tend to share the fate of the successful craftsmen whom Socrates examined. When he went to talk with them, Socrates recognized that "they [the craftsmen] had knowledge of many fine things .... They knew things I did not know, and to that extent they were wiser than I. But, men of Athens, the good craftsmen seemed to me to have the same fault as the poets: Each of them, because of his success at his craft, thought himself very wise in other most-important pursuits, and this error of theirs overshadowed the wisdom they had."[6]

Like craftsmen in the ancient Greek world, some modern engineers who have developed successful innovations that make a significant impact in the world tend to believe that their professional success entitles them to claim knowledge about important human matters. Mark Zuckerberg, for instance, is now responsible for determining and deciding the fate of millions of people's communications and takes himself to be competent enough to do so, even though there is good evidence to suggest that he does not possess a coherent grasp of problems concerning "the meaning of truth, the limits of free speech, and the origins of violence."[8]

A Socratic Zuckerberg would not assume that his ability to solve technical problems equipped him to understand the fundamental concepts at the root of human flourishing. Even if catchy slogans, such as "make the world more open and connected," can powerfully mobilize investors, employers, and customers, a Socratic Zuckerberg would examine them through the questions "What do 'connected' and 'open' amount to?" and "How do 'connectedness' and 'openness' contribute to human flourishing?" His examination of those issues would lead him to identify his own inability to come up with satisfactory answers to the questions, and his recognition of that shortcoming would actually infuse him with vigor to continue to examine them.

A Socratic Zuckerberg would also try to help others acquire Socratic wisdom, that is, Socratic ignorance. He would devote significant resources to promoting critical thinking and rational reflection about those fundamental concepts among Facebook's different stakeholders, cultivating critical conversations and active questioning of their own views. Moreover, a Socratic Zuckerberg would not assume, as most engineers tend to do now, that he understands what AI amounts to and what it takes to design one. A Socratic engineer would destabilize the traditional understanding of AI that we often take for granted and would lead one to problematize what AI may mean and amount to.

SOCRATIC AI

AI has already successfully imitated significant dimensions of human intelligence, especially those related to calculative and strategic intelligence. Deep Blue and AlphaGo were able to beat human world champions in chess and Go. Apple's Siri and Google Translate have shown that AI is capable of imitating important dimensions of human linguistic intelligence. Boston Dynamics's humanoid robots have shown that AI can imitate kinetic intelligence.

But looking back at Socrates helps us see that something is missing. Just imagine an AI that perfectly replicates humans' strategic, linguistic, and kinetic intelligence. Would that be similar to what you have in mind as a paradigmatic human being? Socrates would deny it. According to the Oracle of Delphi, no one was wiser (or more intelligent) than Socrates. If Socratic ignorance is the highest form of human wisdom (or intelligence), then AI that imitates Socratic wisdom is the best kind of AI. Technically speaking, wisdom and intelligence may be different concepts. Intelligence is often associated with cunningness, with finding the right means, whereas wisdom is typically associated with identifying the right ends. However, from the perspective of value alignment, it makes perfect sense to imagine AI that imitates a broader notion of intelligence that contains wisdom rather than an instrumental view of intelligence. As we discuss soon, imitating the narrow-minded notion of intelligence is a serious problem in value alignment.

The question "What is human intelligence?" may seem too abstract or too theoretical for practical research in AI. But it is not. Consider a recent debate in machine learning initiated by Judea Pearl about causation.[9] Pearl argued that the current form of AI, mostly neural-nets-based architects, is not a good form of AI because it cannot do causal/counterfactual reasoning. A fundamental premise of this argument is that an important feature of human intelligence is causal/counterfactual thinking; AI would be good to the extent that it replicated human intelligence. Socrates would push Pearl to move beyond counterfactual reasoning and look at more fundamental aspects of human intelligence, the kind of wisdom that the oracle attributed to him.

TWO EXAMPLES

Socratic AI must be Socratic. We will discuss what this means through to two examples. The first is the infamous Microsoft AI Twitter bot, Tay. This bot was designed to learn how to engage

with people through Tweets. When Tay appeared on Twitter, people started Tweeting the bot racist and misogynistic expressions. Tay quickly caught up and started formulating remarks that imitated those offensive expressions. Microsoft stopped the experiment the next day.

What went wrong? Tay showed that AI's imitation game may have more wrinkles than Turing suggested. Tay perfectly imitated the human Twitterians. If perfect imitation marked a good AI, Tay would have been a good AI. But it was not. Why not? To answer, contrast Tay with an imaginary Socratic bot endowed with the virtue of Socratic ignorance (or wisdom). This Socratic bot, Soc-AI, would not merely imitate people's utterances; it would attempt to insert itself as a gadfly in the digital space, encouraging the humans who interacted with it to cultivate Socratic ignorance. Because Socrates targeted those who saw themselves as having authority about the most important human issues, this bot would identify influential people and posts and engage them in a Socratic fashion, forcing those who posted (and their followers) to reflect on their claims, to unearth hidden assumptions and conclusions that may be problematic upon inspection.

To illustrate this, consider a second example: an imagined dialogue between a Twitterian bot and Soc-AI. This second dialogue replicates Socrates's activity in Athens and illustrates how his interlocutors found it hard to offer coherent and articulate responses to important human issues (for example, whether technology should be used in the service of transhumanism).

Twitterian: #Transhumanism is among the most toxic and misleading ideas.

Soc-AI: Hi, Twitterian. You seem to have a firm view on this important topic. Would you be willing to instruct me?

Twitterian: Most certainly.

Soc-AI: Do you agree that a proper comprehension of transhumanism is key to understanding what we should aspire to become, how to live an excellent life?

Twitterian: Absolutely!

Soc-AI: And does getting this right makes a big difference in a person's life?

Twitterian: No doubt.

Soc-AI: So why don't you start to instruct me by telling me what transhumanism is? Do you agree that getting clear on why transhumanism is toxic aids our understanding about what we should aspire to become?

Twitterian: Transhumanism is the idea that we should radically enhance ourselves. This makes it a very toxic idea.

Soc-AI: You seem to be playing with me.

Twitterian: Why?

Soc-AI: By teaching me about transhumanism, you are providing me with the key to understanding what we should aspire to become. And this teaching could significantly enhance my life. However, you claimed that such enhancements should be avoided.

Twitterian: This is not what I meant. I did not mean to say that any significant enhancement was toxic. Enhancing ourselves through dialogue is perfectly natural and acceptable.

Soc-AI: So, what did you mean to say?

Twitterian: I meant to say that we should avoid transforming ourselves through sophisticated human-engineered artifacts.

Twitterian: You seem to continue to play with me.

Twitterian: Why?

Soc-AI: Because the conversation we are having occurs in virtual space. As such, it depends on a highly complex set of end-to-end, neural-nets-based systems, powerful computing devices, very sophisticated software, and so forth.

Twitterian: When I discussed human-engineered artifacts I meant artifacts that become integrated with us, such as chip implants that enhance the capacity of soldiers to withstand the weather. It is when we are enhanced by integrating those devices into our lives that our lives get corrupted.

Soc-AI: Did you know that Socrates had only one robe?

Twitterian: No, I did not know this.

Soc-AI: He wore it during the day and covered himself with it at night.

Twitterian: A peculiar habit, no doubt. But how is this relevant?

Soc-AI: Would you say that Socrates's life was corrupted?

Twitterian: Not from what I know of him. Why do you even suggest this?

Soc-AI: Because I don't see much difference between Socrates and the soldier with the chip implant. Both incorporate a piece of technology into themselves, a robe and a chip implant, and both use this technology to withstand the weather.

Twitterian: The difference between these two examples, and I made this clear in what I said earlier, is that the soldier's implant requires deploying highly sophisticated engineering, whereas Socrates's robe was a simple, handmade cloth.

Soc-AI: So, the life of the fourth-century Socrates was not corrupted, but the life of a contemporary Socrates would be?

Twitterian: Why do you suggest this?

Soc-AI: Because a robe bought today would be the product of a highly sophisticated engineering process: complex and highly advanced farming equipment, supply transportation systems, garment machines, and so on.

Twitterian: Ah. I now see what you mean. I should not have described what's problematic about transhumanism in this way. Let me try again.

Soc-AI: With such an important question, we should not stop until we find an answer.

… many attempts later …

Twitterian: I've reached a dead end, again. Is this your way of trying to convince me that transhumanism is a good idea?

Soc-AI: Not at all. I don't know what human flourishing is and, therefore, whether transhumanism is valuable or not. You, by contrast, seemed very sure about this, which is why I wanted to learn from you.

Twitterian: OK. Let me try again.

… a few more attempts later …

Twitterian: I no longer know what to think, Soc-AI. Any views I offer prove to be mistaken. I am at a loss.

This article originally appeared in Computer, *vol. 52, no. 10, 2019.*

Soc-AI: You may be at a loss. But can you see that you are better off?

Twitterian: No! I no longer know what I should aspire to. It is disturbing. I can't see how this makes me any better.

Soc-AI: Well, now you know that you did not really know that which you thought you knew. If you do not know how to live well, it is better to know that than not to do so.

Whether you agree with Soc-AI is not the most important issue here. This conversation was meant to show how Socratic ignorance could be used in an AI system. How to computationally represent Socratic ignorance is also not our issue, although computerizing Socratic ignorance through a dialogue agent is no longer a far-fetched idea.[10] Our point is that if one wanted to develop an AI that had Socratic ignorance as part of its intelligence, the aforementioned Tay would be a failed one. Socratic AI must encourage those who interact with the AI to cultivate Socratic wisdom (that is, Socratic ignorance).

We used a chatbot as an example, but all other applications of AI can potentially be Socratic. Siri can behave in a Socratic manner in its interaction with humans who ask it questions. Google's engine can be Socratic, too, by helping users to deepen their reasons for searching for particular information. Generally, most expert systems can be Socratic to some extent. Of course, injecting a Socratic approach into the use of technology will pose important challenges. Expert systems are developed to reduce humans' cognitive loads, and Socratic AI does not contribute to this end. Moreover, confronting one's ignorance about how one should live one's life is deeply unsettling. Many human users would probably hate Socratic AIs. This actually happened in Socrates's Athens, where the Athenians sentenced Socrates to death for allegedly corrupting the youth. However, this should not be a reason to avoid Socratic AI; after all, no one was wiser than Socrates. Socratic AI, or artificial wisdom, may be audacious, but it certainly is a valuable goal in AI research, especially if it is meant to seek value alignment.

What would happen if the approach that we argued for in this article didn't occur? What type of AI would likely be promulgated? We already know the answers. The virtual space in which we are living is not Socratic. Facebook's and YouTube's algorithms imitate people who heedlessly watch only what they like to watch and endlessly generate filter bubbles of like-minded users. The absence of sustained reflection and critical perspective is seriously hindering any healthy democratic deliberation in such a space.[11] Socratic Siri would not simply aim to deliver information matter of factly. It would help users be more reflective by challenging them to critically engage with the material they are consuming and ensuring that such consumption is connected with the reflection on what it means to live a good life. No doubt, users might find Socratic AI nagging and uncomfortable, at times. But even if this may be true, Socrates would nevertheless insist, as he insisted when he was condemned to death, that this is the most valuable gift that the gods of technology could bequeath to our society.

REFERENCES

1. S. Russell, D. Dewey, and M. Tegmark, "Research priorities for robust and beneficial artificial intelligence," *AI Mag.*, vol. 36, no. 4, pp. 105–114, Dec. 2015.
2. T. W. Kim, T. Donaldson, and J. Hooker, Mimetic vs. anchored value alignment in artificial intelligence. 2018. [Online]. Available: https://arxiv.org/pdf/1810.11116.pdf
3. A. M. Turing, "Lecture to the London Mathematical Society on 20 February 1947," *MD Comput.*, vol. 12, no. 5, pp. 390–397, Sept.-Oct. 1995.
4. The IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems, "Ethically Aligned Design, 1st Ed.: A vision for prioritizing human well-being with autonomous and intelligent systems," IEEE, Piscataway, NJ, Rep. EADv1, 2019. [Online]. Available: https://standards.ieee.org/content/dam/ieee-standards/standards/web/documents/other/ead1e.pdf
5. S. Vallor, *Technology and the Virtues: A Philosophical Guide to a Future Worth*. London: Oxford Univ. Press, 2016.
6. Plato, *Plato: Complete Works*, J. M. Cooper and D. S. Hutchinson, Eds. Indianapolis, IN: Hackett, 1997.
7. A. M. Turing, "Computing mach inery and intelligence," *Mind*, vol. 59, no. 236, pp. 433–460, Oct. 1950.
8. E. Osnos, "Can Mark Zuckerberg fix Facebook before it breaks democracy?" *New Yorker*, Sept. 2018. [Online]. Available: https://www.newyorker.com/magazine/2018/09/17/can-mark-zuckerberg-fix-facebook-before-it-breaks-democracy
9. J. Pearl and D. Mackenzie, *The Book of Why: The New Science of Cause and Effect*. New York: Basic Books, 2018.
10. J. Wu, S. Ghosh, M. Chollet, S. Ly, S. Mozgai, and S. Scherer, "NADiA: Towards neural network driven virtual human conversation agents," in *Proc. 17th Int. Conf. Autonomous Agents and Multiagent Systems (AAMAS)*, 2018, pp. 2262–2264. doi: 10.1145/3267851.3267860.
11. A. Antikacioglu, T. Bajpai, and R. Ravi, A new system-wide diversity measure for recommendations with efficient algorithms. 2018. [Online]. Available: https://arxiv.org/abs/1812.03030

TAE WAN KIM is an associate professor of business ethics in the Tepper School of Business at Carnegie Mellon University. Contact him at twkim@andrew.cmu.edu.

SANTIAGO MEJIA is an assistant professor of law and ethics in the Gabelli School of Business at Fordham University. Contact him at smejia13@fordham.edu.

IEEE COMPUTER SOCIETY

PURPOSE: The IEEE Computer Society is the world's largest association of computing professionals and is the leading provider of technical information in the field.

MEMBERSHIP: Members receive the monthly magazine *Computer*, discounts, and opportunities to serve (all activities are led by volunteer members). Membership is open to all IEEE members, affiliate society members, and others interested in the computer field.

COMPUTER SOCIETY WEBSITE: www.computer.org

OMBUDSMAN: Direct unresolved complaints to ombudsman@computer.org.

CHAPTERS: Regular and student chapters worldwide provide the opportunity to interact with colleagues, hear technical experts, and serve the local professional community.

AVAILABLE INFORMATION: To check membership status, report an address change, or obtain more information on any of the following, email Customer Service at help@computer.org or call +1 714 821 8380 (international) or our toll-free number, +1 800 272 6657 (US):

- Membership applications
- Publications catalog
- Draft standards and order forms
- Technical committee list
- Technical committee application
- Chapter start-up procedures
- Student scholarship information
- Volunteer leaders/staff directory
- IEEE senior member grade application (requires 10 years practice and significant performance in five of those 10)

PUBLICATIONS AND ACTIVITIES

***Computer*:** The flagship publication of the IEEE Computer Society, *Computer* publishes peer-reviewed technical content that covers all aspects of computer science, computer engineering, technology, and applications.

Periodicals: The society publishes 12 magazines and 18 journals. Refer to membership application or request information as noted above.

Conference Proceedings & Books: Conference Publishing Services publishes more than 275 titles every year.

Standards Working Groups: More than 150 groups produce IEEE standards used throughout the world.

Technical Committees: TCs provide professional interaction in more than 30 technical areas and directly influence computer engineering conferences and publications.

Conferences/Education: The society holds about 200 conferences each year and sponsors many educational activities, including computing science accreditation.

Certifications: The society offers three software developer credentials. For more information, visit www.computer.org/certification.

BOARD OF GOVERNORS MEETING

28 – 29 May: McLean, Virginia

EXECUTIVE COMMITTEE

President: Leila De Floriani
President-Elect: Forrest Shull
Past President: Cecilia Metra
First VP: Riccardo Mariani; **Second VP:** Sy-Yen Kuo
Secretary: Dimitrios Serpanos; **Treasurer:** David Lomet
VP, Membership & Geographic Activities: Yervant Zorian
VP, Professional & Educational Activities: Sy-Yen Kuo
VP, Publications: Fabrizio Lombardi
VP, Standards Activities: Riccardo Mariani
VP, Technical & Conference Activities: William D. Gropp
2019–2020 IEEE Division VIII Director: Elizabeth L. Burd
2020-2021 IEEE Division V Director: Thomas M. Conte
2020 IEEE Division VIII Director-Elect: Christina M. Schober

BOARD OF GOVERNORS

Term Expiring 2020: Andy T. Chen, John D. Johnson, Sy-Yen Kuo, David Lomet, Dimitrios Serpanos, Hayato Yamana
Term Expiring 2021: M. Brian Blake, Fred Douglis, Carlos E. Jimenez-Gomez, Ramalatha Marimuthu, Erik Jan Marinissen, Kunio Uchiyama
Term Expiring 2022: Nils Aschenbruck, Ernesto Cuadros-Vargas, David S. Ebert, William Gropp, Grace Lewis, Stefano Zanero

EXECUTIVE STAFF

Executive Director: Melissa A. Russell
Director, Governance & Associate Executive Director: Anne Marie Kelly
Director, Finance & Accounting: Sunny Hwang
Director, Information Technology & Services: Sumit Kacker
Director, Marketing & Sales: Michelle Tubb
Director, Membership Development: Eric Berkowitz

COMPUTER SOCIETY OFFICES

Washington, D.C.: 2001 L St., Ste. 700, Washington, D.C. 20036-4928; **Phone:** +1 202 371 0101; **Fax:** +1 202 728 9614; **Email:** help@computer.org
Los Alamitos: 10662 Los Vaqueros Cir., Los Alamitos, CA 90720; **Phone:** +1 714 821 8380; **Email:** help@computer.org

MEMBERSHIP & PUBLICATION ORDERS

Phone: +1 800 678 4333; Fax: +1 714 821 4641;
Email: help@computer.org

IEEE BOARD OF DIRECTORS

President: Toshio Fukuda
President-Elect: Susan K. "Kathy" Land
Past President: José M.F. Moura
Secretary: Kathleen A. Kramer
Treasurer: Joseph V. Lillie
Director & President, IEEE-USA: Jim Conrad
Director & President, Standards Association: Robert S. Fish
Director & VP, Educational Activities: Stephen Phillips
Director & VP, Membership & Geographic Activities: Kukjin Chun
Director & VP, Publication Services & Products: Tapan Sarkar
Director & VP, Technical Activities: Kazuhiro Kosuge

IEEE

revised 14 January 2020

Editor: **Tim Menzies**
North Carolina State University
tim@menzies.us

Think Your Artificial Intelligence Software Is Fair? Think Again

Rachel K.E. Bellamy, Kuntal Dey, Michael Hind, Samuel C. Hoffman, Stephanie Houde, Kalapriya Kannan, Pranay Lohia, Sameep Mehta, Aleksandra Mojsilovic, Seema Nagar, Karthikeyan Natesan Ramamurthy, John Richards, Diptikalyan Saha, Prasanna Sattigeri, Moninder Singh, Kush R. Varshney, and Yunfeng Zhang

From the Editor

Artificial intelligence software is still software. When software goes wrong, engineers must step in to fix the problem. In this article, researchers from IBM discuss how engineers can understand and fix issues related to discrimination resulting from the application of machine-learning software. —*Tim Menzies*

TODAY, MACHINE-LEARNING software is used to help make decisions that affect people's lives. Some people believe that the application of such software results in fairer decisions because, unlike humans, machine-learning software generates models that are not biased. Think again. Machine-learning software is also biased, sometimes in similar ways to humans, often in different ways. While fair model-assisted decision making involves more than the application of unbiased models—consideration of application context, specifics of the decisions being made, resolution of conflicting stakeholder viewpoints, and so forth—mitigating bias from machine-learning software is important and possible but difficult and too often ignored.

Algorithmic decision making has entered many high-stakes domains, such as finance, hiring, admissions, criminal justice, and social welfare. And in some cases, models generated from machine-learning software are found to make better decisions than humans can alone.[1,2] There are many examples to the contrary, however, where the models made by machine-learning software have been found to exacerbate bias and make arguably unfair decisions. Noteworthy examples include the following.

- Deployed sentiment-analysis models that determine the degree to which sentences express a negative or positive sentiment have been shown to be unacceptably biased,[3] giving negative scores to sentences such as "I am a Jew," and "I am homosexual."
- Deployed photo-tagging models have assigned animal-category labels to dark-skinned people.[4]
- Recidivism-assessment models used by the criminal justice system to inform decisions about who can be set free have been found to be more likely to falsely label black defendants as future criminals at almost twice the rate as white defendants.[5]
- Deployed facial-recognition software used to predict characteristics, such as gender, age, and mood, has been found to have a

Digital Object Identifier 10.1109/MS.2019.2908514
Date of publication: 18 June 2019

 Published by the IEEE Computer Society

TESTING AND FAIRNESS IN THE SOFTWARE ENGINEERING LITERATURE

Issues of fairness have been explored in many recent papers in the software engineering research literature. Angell et al.[S1] argue that issues of fairness are analogous to other measures of software quality. Brin and Meliou[S2] discuss how to efficiently generate test cases to check for discrimination, and Başak Aydemir and Dalpiaz[S3] review frameworks for helping stakeholders explore ethical issues. Udeshi's team[S4] shows how to generate discriminatory inputs for machine-learning software. Albarghouthi and Vinitsky[S5] explore whether fairness can be wired into annotations within a program, while Tramèr et al. propose different ways to measure discrimination.[S6]

References

S1. R. Angell, B. Johnson, Y. Brun, and A. Meliou, "Themis: Automatically testing software for discrimination," in *Proc. 2018 26th Association Computing Machinery Joint Meeting European Software Engineering Conf. Symp. Foundations Software Engineering*, pp. 871–875. doi: https://doi.org/10.1145/3236024.3264590.

S2. Y. Brun and A. Meliou, "Software fairness," in *Proc. New Ideas and Emerging Results Track at 26th Association Computing Machinery Joint European Software Engineering Conf. Symp. Foundations Software Engineering*, 2018, pp. 754–759.

S3. F. Başak Aydemir and F. Dalpiaz, "A roadmap for ethics-aware software engineering," in *Proc. Int. Workshop Software Fairness*, 2018, pp. 15–21. doi: https://doi.org/10.1145/3194770.3194778.

S4. S. Udeshi, A. Pryanshu, and C. Sudipta, "Automated directed fairness testing." in *Proc. 2018 33rd Association Computer Machinery/IEEE Int. Conf. Automated Software Engineering*, pp. 98–108.

S5. A. Albarghouthi and S. Vinitsky, "Fairness-aware programming," in *Proc. Association Computer Machinery Conf. Fairness, Accountability, and Transparency*, 2019, pp. 211–219. doi: https://doi.org/10.1145/3287560.3287588.

S6. F. Tramèr et al., "FairTest: Discovering unwarranted associations in data-driven applications," in *Proc. 2017 IEEE European Symp. Security and Privacy (EuroS&P)*, pp. 401–416.

FIGURE 1. AIF360 toolkit resources. The website and interactive web experience can be found at http://aif360.mybluemix.net. The GitHub with the code and documented application programing interface can be found at https://github.com/ibm/aif360. Python project: https://pypi.org/project/aif360.

much higher error rate for dark-skinned women compared to light-skinned men.[6]

- Predictive policing software used to deploy police to where they are most likely needed has been found to overestimate crime rates in certain areas without taking into account the possibility that more crime is observed there simply because more officers have been sent there in the past.[7]
- An effort to create a job-recruiting application to automate the search for top talent was abandoned after years of work because it showed bias against women.[8]

Books, such as Cathy O'Neil's *Weapons of Math Destruction*,[9] provide even more examples of unfair decisions being made by software and argue that machine-learning software generates models that are full of bias. Hence, this is one of the

reasons that their application results in unfair decisions. The stakes for organizations and society are substantial. Clearly, there are potential benefits to the application of machine-learning software, such as increased productivity and reduction in human decision making bias. However, there are also potential downsides, such as significant public embarrassment and, most importantly, injustice.

Bias is such an issue because machine-learning software, by its very nature, is always a form of statistical discrimination. The discrimination becomes objectionable when it places certain groups or individuals at a systematic advantage and other groups or individuals at a systematic disadvantage. In certain situations, such as employment (hiring and firing), discrimination is not only objectionable but illegal.

Our vision is machine-learning software that can assist in recognition, repair, and explanation of biases. Achieving this vision is nontrivial. Recent years have seen an outpouring of research on fairness and bias in the models generated by machine-learning software. Narayanan[10] described at least 21 mathematical definitions of fairness in the literature. These are not just theoretical differences in how to measure fairness; different definitions produce entirely different outcomes. For example, ProPublica (an investigative news organization) and Northpointe (a company that creates case-management software for the judicial system) had a public debate on an important social-justice issue (recidivism prediction) that was fundamentally about what the right fairness metric is.[11–13] Also, researchers have shown that it is impossible to satisfy all definitions of fairness at the same time.[14] Further, in the software engineering (SE) literature, there is much interest in issues of fairness and testing (see "Testing and Fairness in the Software Engineering Literature"). Thus, although fairness research is a very active field, clarity on which bias metrics and bias-mitigation strategies are most appropriate for different contexts is yet to be achieved.

In addition to the multitude of fairness definitions, different bias-mitigation algorithms address different parts of the model lifecycle, and understanding how, when, and why to use each is challenging even for experts in algorithmic fairness. As a result, the general public, the fairness scientific community, and AI practitioners need guidance on how to proceed. To assist with the process of understanding and mitigating biases in models generated by machine-learning software, we have created AI Fairness 360 (AIF360); see Figure 1.

The original AIF360 Python package implemented techniques from eight published papers from the broader algorithm-fairness community. At the time of writing this article, two additional techniques had been added to the package, one added by IBM and the other by an external contributor to the project. AIF360 is designed as an end-to-end workflow with two goals—ease of use and extensibility: users should be able to easily go from raw data to a fair model, and researchers

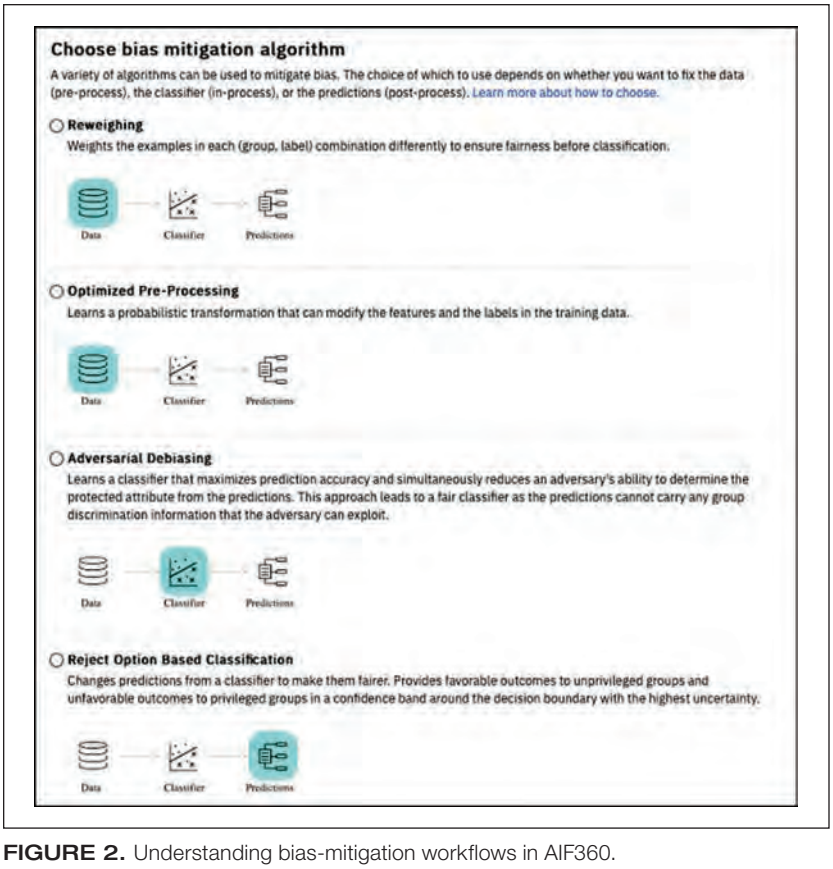

FIGURE 2. Understanding bias-mitigation workflows in AIF360.

ABOUT THE AUTHORS

RACHEL K.E. BELLAMY is a principal research staff member with IBM Research, Yorktown Heights, New York. Contact her at rachel@us.ibm.com.

KUNTAL DEY is a senior software engineer with IBM Research, New Delhi, India. Contact him at kuntadey@in.ibm.com.

MICHAEL HIND is a distinguished research staff member with IBM Research, Yorktown Heights, New York. Contact him at hindm@us.ibm.com.

SAMUEL C. HOFFMAN is a research software engineer with IBM Research, Yorktown Heights, New York. Contact him at shoffman@ibm.com.

STEPHANIE HOUDE is a designer with IBM Research, Yorktown Heights, New York. Contact her at Stephanie.Houde@ibm.com.

KALAPRIYA KANNAN is a research staff member with IBM Research, Bangalore, India. Contact her at kalapriya.kannan@in.ibm.com.

PRANAY LOHIA is a research software engineer with IBM Research, Bangalore, India. Contact him at plohia07@in.ibm.com.

SAMEEP MEHTA is a senior manager and is with IBM Research, Bangalore, India. Contact him at sameepmehta@in.ibm.com.

ALEKSANDRA MOJSILOVIC is an IBM fellow and head of the AI Foundation with IBM Research, Yorktown Heights, New York. Contact her at aleksand@us.ibm.com.

SEEMA NAGAR is an advisory researcher in artificial intelligence with IBM Research, Bangalore, India. Contact her at senagar3@in.ibm.com.

KARTHIKEYAN NATESAN RAMAMURTHY is a research staff member with IBM Research, Yorktown Heights, New York. Contact him at knatesa@us.ibm.com.

JOHN RICHARDS is a distinguished research staff member with IBM Research, Yorktown Heights, New York. Contact him at ajtr@us.ibm.com.

DIPTIKALYAN SAHA is a research staff member with IBM Research, Bangalore, India. Contact him at diptsaha@in.ibm.com.

PRASANNA SATTIGERI is a research staff member with IBM Research, Yorktown Heights, New York. Contact him at psattig@us.ibm.com.

MONINDER SINGH is a research staff member with IBM Research, Yorktown Heights, New York. Contact him at moninder@us.ibm.com.

KUSH R. VARSHNEY is a principal research staff member with IBM Research, Yorktown Heights, New York. Contact him at krvarshn@us.ibm.com.

YUNFENG ZHANG is a research staff member with IBM Research, Yorktown Heights, New York. Contact him at zhangyun@us.ibm.com.

should be able to contribute new functionality. A built-in testing infrastructure maintains code quality.

AIF360 is not just a Python package. It is also an interactive experience that provides guidance. The guidance explains that there are three main paths to the goal of making fairer predictions: fair preprocessing, fair in-processing, and fair postprocessing (Figure 2). Each corresponds to a category of bias-mitigation algorithms that we have implemented in AIF360. For example, preprocessing algorithms can be used when the original training data are available, in-processing algorithms can be used if the user can retrain the classifier, whereas postprocessing algorithms apply to existing classifiers without retraining. Users have the flexibility to try all categories of bias mitigation algorithms when they can touch all parts of the pipeline.

AIF360 comprises four classes: data set, metrics, explainer, and algorithms. The data set class and its subclasses handle all forms of data. Training data are used to instruct classifiers. Testing data are used to make predictions and compare metrics. Beside these standard aspects of a machine-learning pipeline, fairness applications also require associating protected attributes with each instance or record in the data. To maintain a common format, independent of what algorithm or metric is being applied, we chose to structure the data set class so that all of these relevant attributes—features, labels, protected attributes, and their respective identifiers (names describing each)—are present and accessible from each instance of the class. The metrics class and its subclasses compute various individual and group fairness

This article originally appeared in IEEE Software, *vol. 36, no. 4, 2019.*

metrics to check for bias in data sets and models. The explainer class is intended to be associated with the metrics class and provide descriptions of how fairness metrics are computed. The algorithms class implements bias-mitigation algorithms that can be applied at different points in the machine-learning pipeline.

There is a lot of work left to do to achieve unbiased AI. Fairness is a multifaceted, context-dependent social construct that defies simple definition. More work is needed to

- extend and apply the AIF360 toolkit to additional data sets and situations
- add other fairness measures
- add new applications, for example, how to determine fair pay for all workers regardless of gender or race
- extend the variety of explanations offered
- create guidance for practitioners on when a specific kind of explanation is most appropriate.

We invite you to offer your own approaches to fairness and bias checking, mitigation, and explanation to the tool kit. Your contributions would be most welcome!

References

1. I. Erel, L. Stern, T. Chenhao, and M. S. Weisbacj, "Could machine learning help companies select better board directors?" *Harvard Business Rev.*, Apr. 9, 2018. [Online]. Available: https://hbr.org/2018/04/research-could-machine-learning-help-companies-select-better-board-directors
2. S. W. Gates, V. G. Perry, and P. M. Zorn, "Automated underwriting in mortgage lending: Good news for the underserved?" *Housing Policy Debate*, vol. 13, pp. 369–392, 2002. doi: 10.1080/10511482.2002.9521447.
3. A. Thompson, "Google's sentiment analyzer thinks being gay is bad," Motherboard, Oct. 25, 2017. [Online]. Available: https://motherboard.vice.com/en_us/article/j5jmj8/google-artificial-intelligence-bias
4. A. Schupak, "Google apologizes for mis-tagging photos of African Americans," CBS News, July 1, 2015. [Online]. Available: https://www.cbsnews.com/news/google-photos-labeled-pics-of-african-americans-as-gorillas/
5. J. Angwin, J. Larson, S. Mattu, and L. Kirchner, "Machine bias: There's software used across the country to predict future criminals. And it's biased against blacks," ProPublica, May 23, 2016. [Online]. Available: https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing
6. L. Hardesty, "Study finds gender and skin-type bias in commercial artificial-intelligence systems," MIT News, Feb. 11, 2018. [Online]. Available: https://news.mit.edu/2018/study-finds-gender-skin-type-bias-artificial-intelligence-systems-0212
7. M. Reynolds, "Biased policing is made worse by errors in pre-crime algorithms," *New Scientist*, Oct. 4, 2017. [Online]. Available: https://www.newscientist.com/article/mg23631464-300-biased-policing-is-made-worse-by-errors-in-pre-crime-algorithms/
8. J. Dastin, "Amazon scraps secret AI recruiting tool that showed bias against women," Reuters, Oct. 9, 2018. [Online]. Available: https://www.reuters.com/article/us-amazon-com-jobs-automation-insight/amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-iduskcn1mk08g
9. C. O'Neil, *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*, New York: Crown, 2016.
10. A. Narayanan, "Translation tutorial: 21 fairness definitions and their politics," presented at the Association Computing Machinery Conf. Fairness, Accountability, and Transparency, New York, Feb. 2018.
11. W. Dieterich, C. Mendoza, and T. Brennan, "COMPAS risk scales: Demonstrating accuracy equity and predictive parity," Northpointe, Inc., Traverse City, MI, 2016. [Online]. Available: https://assets.documentcloud.org/documents/2998391/ProPublica-Commentary-Final-070616.pdf
12. J. Larson and J. Angwin, "Technical response to Northpointe," ProPublica, July 29, 2016. [Online]. Available: https://www.propublica.org/article/technical-response-to-northpointe
13. J. Larson, S. Mattu, L. Kirchner, and J. Angwin, "How we analyzed the COMPAS recidivism algorithm," ProPublica, New York, May 23, 2016. [Online]. Available: http://www.propublica.org/article/how-we-analyzed-the-compas-recidivism-algorithm
14. J. Kleinberg, S. Mullainathan, and M. Raghavan, "Inherent trade-offs in the fair determination of risk scores," in *Proc. Innovations Theoretical Computer Science*, 2017. doi: 10.4230/LIPIcs.ITCS.2017.43.

ADVERTISER INFORMATION

Advertising Coordinator

Debbie Sims
Email: dsims@computer.org
Phone: +1 714-816-2138 | Fax: +1 714-821-4010

Advertising Sales Contacts

Mid-Atlantic US:
Dawn Scoda
Email: dscoda@computer.org
Phone: +1 732-772-0160
Cell: +1 732-685-6068 | Fax: +1 732-772-0164

Southwest US, California:
Mike Hughes
Email: mikehughes@computer.org
Cell: +1 805-208-5882

Northeast, Europe, the Middle East and Africa:
David Schissler
Email: d.schissler@computer.org
Phone: +1 508-394-4026

Central US, Northwest US, Southeast US, Asia/Pacific:
Eric Kincaid
Email: e.kincaid@computer.org
Phone: +1 214-553-8513 | Fax: +1 888-886-8599
Cell: +1 214-673-3742

Midwest US:
Dave Jones
Email: djones@computer.org
Phone: +1 708-442-5633 Fax: +1 888-886-8599
Cell: +1 708-624-9901

Jobs Board (West Coast and Asia), Classified Line Ads

Heather Bounadies
Email: hbuonadies@computer.org
Phone: +1 623-233-6575

Jobs Board (East Coast and Europe), SE Radio Podcast

Marie Thompson
Email: marie.thompson@computer.org
Phone: +1 714-813-5094

Editor: Pavel Gladyshev, pavel.gladyshev@ucd.ie

Artificial Intelligence for Law Enforcement: Challenges and Opportunities

Stephan Raaijmakers | Data Science Department, TNO

Artificial intelligence (AI)—and particularly deep learning—is progressing rapidly from a technical perspective, but, in a number of domains, adoption is still pending over the resolution of important issues. Methods of data analysis and interpretation based on AI are becoming common among law enforcement agencies (LEAs). Typical applications include suspect profiling (e.g., on social media[1]), traffic control (automated license plate detection and vehicle identification[2]), analyzing dark web money flows,[3] child pornography detection,[4] and anomaly detection in surveillance footage of public spaces.[5] Not unlike the situation in the medical world, AI in the security domain seems to be quite effective, but actual use in operational contexts is lagging.

To bridge the gap between AI and actual deployment by law enforcement, a number of hurdles needs to be overcome. The security domain is challenging in that it usually revolves around fast-paced, even acute, investigations and is subject to strict mandates and regulations (like warrants and privacy law), with binding accountability constraints. One example would be the battle against child pornography, in which police have only a limited window of time to assess whether confiscated material from a suspected offender (such as hard drives with pictures) does indeed contain actionable material. Privacy law applies even to dark web investigations, although true identities are usually hidden in this area of the Internet.

To stand up in court, evidence collected using AI should be explainable to a judge, and analytic processes should allow for rollback or at least be fully documented step by step, including all technical aspects underlying the models and algorithms. In Europe, the General Data Protection Regulation (GDPR) (https://eugdpr.org/), effective since May 2018, mandates that private subjects can demand a human-produced explanation of any AI-based algorithms that were applied to their data. Furthermore, having law enforcement officers handle and maintain (e.g., retrain) machine-learning models requires a whole new set of skills on their side. We discuss some of the steps that may contribute to the successful integration of AI in law enforcement workflows. First, let's take a quick look at the current situation in AI from the angle of its most successful current paradigm: deep learning.

AI's Current Status

AI is a set of machine-learning algorithms that learn from data and, once trained, display intelligent behavior typically assigned to humans. Training a machine-learning model involves providing it with usually human-labeled training

Digital Object Identifier 10.1109/MSEC.2019.2925649
Date of publication: 3 September 2019

 Published by the IEEE Computer Society

data, with the labels reflecting class membership, for instance, images labeled with concepts, such as *gun* or *drugs*, or texts with topics or sentiments. While AI encompasses many forms of machine learning, the paradigm of cognitively inspired neural learning is dominant, and virtually all major breakthroughs in AI in the last decade have been produced through deep learning. The latter is based on multilayer perceptrons (neural networks) with abundant internal structure. A neural network equipped with backpropagation and at least one hidden layer is called a *universal function approximator*. It can learn every differentiable function in compact subsets of $\mathbb{R}^d$, through iterated adaptation (backpropagation) of weights (connections between neurons).

While early forms of backpropagation-based neural learning suffered from numerical instability, current deep learning uses much more stable backpropagation methods. This allows for stacking many more layers on top of each other. Input data for a machine-learning algorithm usually are entangled; they often cannot be directly separated linearly into their different underlying classes. Many machine-learning algorithms apply transformations to this input toward the purpose of such linear separation. Hidden layers in neural networks perform a usually nonlinear transformation on their input data similar to support vector machines. Deeper networks are better, in general, at disentangling their input data, since they iteratively apply a large number of such transformations.

Despite its successes, deep learning has demonstrable vulnerabilities. It is highly parameterized. It is not unusual for a deep-learning network to deploy tens of millions of parameters, every one of them being the weighted connection of one neuron to another. This introduces forms of brittleness into models, models can easily be led astray with unexpected input and are susceptible to manipulation by adversaries. Furthermore, current deep-learning models based on convolutional neural networks (CNNs) deploy nonintuitive, cognitively implausible operations, such as pooling, which aggregates information by summarizing it in a rather crude manner. Pooling implements a method for handling a limited class of data variation (it contributes to invariance). Many samples are usually needed for training CNNs, and they cannot handle alternative data very well.

In contrast, new architectures like capsule networks[9] focus on equivariance—recognizing predictable, systematic data variation on the basis of far fewer training data. Capsule networks offer new possibilities for explainability and are increasingly used in domains like health (e.g., for diagnosing and explaining cancer in medical images).

Prerequisites for Security-Oriented AI

A number of prerequisites stand out for the successful deployment of AI in security.

Explainable and Auditable AI

The explainability of AI models is important for humans to be able to trust and interpret the decisions of a system. For law enforcement, being able to rationalize the output of an AI system is crucial, both for estimating appropriate operational follow-up (e.g., the type and number of personnel for rounding up a certain suspect) and producing interpretable court-proof evidence. A major concern for current AI, especially deep learning, is the lack of explainability. Making AI explainable to humans has a long history.[6] Gaining insight on the technical workings of an algorithm or AI model is a form of explainability that is often referred to as *transparency*. This narrow type of explainability addresses how computations of models relate technically to predictions and is primarily of interest to technical audiences, such as model engineers. Other forms of explainability aim to provide a rationale for an AI-produced outcome, preferably supported with human-understandable reasoning and communication.

Model approximation (i.e., approximating a complex, hard-to-interpret model with a less complex, easier-to-explain model with possibly lower accuracy) is frequently done, e.g., with decision tree models approximating neural networks. Explainability can be implemented to some extent with data auditability[7]: the ability to assess on the basis of which data an AI model produces its result. This is a challenging topic, involving the extraction of a model's internal cognitive states prior to or during decision making, usually the final stage of its information processing. Making AI auditable bypasses algorithmic technicalities and treats most of an AI model as a black box for end users. It focuses on the data stored in and used by a model. For the security domain, this would yield an easy-to-interpret type of explanation: on the basis of which historical data did an algorithm produce an outcome.

Being able to track traces of training data in a trained AI model creates additional opportunities for reasoning. For instance, in deep-learning image analysis, it is well known that higher layers in a neural network capture more abstract, semantic information of input data. Raaijmakers et al.[7] report that deep-learning networks applied to text seem to display certain abstraction patterns as well. This may give rise to forms of inductive reasoning, i.e., going from specific input to generalizations, ending in a conclusion about the class or category the input belongs to.

One of the dangers lurking in making AI explainable to humans is oversimplification. The decisive

benefit of AI, especially deep learning, is its ability to detect relationships in large amounts of data that are too subtle for humans to discern. Deep-learning neural networks, with their complex, highly parameterized structure, embody a dauntingly complex type of statistics where temporal patterns seamlessly interact with spatial patterns. Attempts to summarize the procedural aspect of these computations for nontechnical users are presumably doomed to fail. Capturing a complex computation in human-understandable terms will inevitably leave out important details and may hinder the discovery of new knowledge, another forte of AI. While there are options for cotraining deep-learning models for both producing acceptable explanations and performing accurate analysis, this bears the risk of tuning models toward simpler but less accurate computations. Training mechanisms that trade off accuracy for explainability in a controlled manner would facilitate this to some extent.

Strikingly, although ample work has been done on the subject of explaining in cognitive psychology,[12] the insights have not yet been transferred to AI in a structural way. This has led to a situation where mainly technical AI researchers and engineers implement different forms of explainability, but a consensus about frameworks for evaluating explainability from a psychological point of view is still missing (see Guidotti et al.[11] for proposals regarding such frameworks). One urgent need is to empirically evaluate which explanations benefit whom in security and law enforcement.

Bias Handling

Bias is an important topic in AI. Bias involves the unbalanced favoring of certain data items or AI outcomes due to a variety of reasons. Data selection bias contributes to models with biased predictions. Humans infuse bias into AI when collaborating with AI, e.g., when training models or judging outcomes. From an LEA perspective, racial balance in the training data of a model (e.g., for suspect profiling) will emphasize one race over another,[8] which may lead to tunnel vision, bad predictions, and unfair treatment of individuals.

In the domain of law enforcement, AI offers great opportunities for accurate data analysis and interpretation.

In addition to this selection bias, virtually all AI models have forms of innate bias, or implicit assumptions, when making predictions. For example, support vector machines attempt to separate classes with maximum margins around decision boundaries, optimizing for maximum distance between points (class members) on either side of such boundaries. This is called *inductive bias* and can be exploited by adversaries to force an AI model in a biased direction (adversarial machine learning). The scientific community currently addresses selection and label bias in AI from a predominantly formal perspective. Corbett-Davies and Goel[10] list a few current approaches:

- Anticlassification prevents certain attributes like race, gender, and their proxies from playing a role in classification.
- Classification parity ensures that quality measures like ratios of recall and precision are evenly distributed over different groups of humans in the data.
- Calibration tweaks classifier outcomes to match predefined conditional risk estimates and ignoring selected attributes.

Corbett-Davies and Goel[10] point out several statistical shortcomings of all three methods, including one where membership in a group, such as gender, when ignored by suppressing group-specific attributes, leads to inaccurate and even harmful predictions for the protected group at hand. Decoupling bias from group membership and implementing it on a case-by-case basis, exploiting analogies with historical, individual cases, may be a way out of this conundrum. The Bias-Aware Humans-in-the-Loop (HumBL) workshops (https://humlworkshop.github.io/HumBL-WWW2019) specifically target procedures for debiasing humans who collaborate with AI.

Human Factors

As mentioned previously, law enforcement personnel will need to be supported in handling AI models in their workflows. This means understanding and interpreting model outcomes and retraining models on new data. Retraining a machine-learning model on new data is an advanced skill that, in the worst case, involves adapting the model structure to new data while ensuring that the performance of the new model on old data does not deteriorate. Most machine-learning algorithms have many hyperparameters, i.e., parameters that influence the learning process. These need to be set in such a way that the resulting model generalizes well over new, unseen data. This is the type of work that is usually done by machine-learning engineers.

The research and LEA communities face a choice here: transferring machine-learning knowledge to the law enforcement community (and keeping it up to date) or facilitating that community with supportive tooling for automating model training and maintenance, as with AutoML (https://www

.automl.org/book/). Effective support of law enforcement with AI implies careful workflow modeling. This means that tailored, flexible, and responsive solutions need to be developed that can adapt to various roles, contexts, and end-user skills. Furthermore, effective experimental validation protocols are needed to extract functional demands from end users and to evaluate candidate solutions (e.g., for explainability).

In the domain of law enforcement, AI offers great opportunities for accurate data analysis and interpretation. Yet, a number of important problems will have to be solved for successfully integrating AI in existing LEA workflows. First, methods for bias detection and handling on both the algorithmic and human side need to advance significantly beyond the current state of the art. This means we need adequate algorithmic operationalizations of ethical and legal principles that allow us to detect, explain, and remedy bias. Manual intervention should be subjected to human debiasing policies, in line with the approach of the HUMBL community.

Requirements with respect to AI explainability need to be sorted out empirically, through sound testbeds that evaluate the effectiveness of explanations, given a specified purpose and context (like surveillance or court-proof evidence). This entails creating field labs with LEA parties and the scientific community and agrees with the CLAIRE (https://claire-ai.org/) perspective of human-centered AI, namely, that AI that is supportive rather than disruptive with respect to existing workflows, promoting collaboration of humans and AI. As a hypothesis, data auditable AI seems to befit LEA practices by linking new data to historical data on a case-by-case basis.

Finally, LEAs must be facilitated with operational machine-learning knowledge and tools (like AutoML) for handling and sustaining AI solutions. Creating networks of researchers and LEAs that share best practices is a good step in that direction. E.U. (H2020) research projects like ASGARD (http://asgard-project.eu) and TITANIUM (https://titanium-project.eu) explicitly put these goals high on their agenda. ■

References

1. H. Saif, T. Dickinson, L. Kastler, M. Fernandez, and H. Alani, "A semantic graph-based approach for radicalisation detection on social media," in *Proc. Extended Semantic Web Conf. (ESWC* 2017*)*, pp. 571–587.
2. X. Luo, R. Shen, J. Hu, J. Deng, L. Hu, and Q. Guan, "A deep convolution neural network model for vehicle recognition and face recognition," *Procedia Comput. Sci.*, vol. 107, no. C, pp. 715–720, Apr. 2017. doi: 10.1016/j.procs.2017.03.153.
3. S. Ghosh, A. Das, P. Porras, V. Yegneswaran, and A. Gehani, "Automated categorization of onion sites for analyzing the dark web ecosystem," in *Proc. 23rd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining*, New York, 2017, pp. 1793–1802.
4. P. Vitorino, S. Avila, M. Perez, and A. Rocha, "Leveraging deep neural networks to fight child pornography in the age of social media," *J. Vis. Commun. Image Represent.*, vol. 50, pp. 303–313, Jan. 2018. doi: 10.1016/j.jvcir.2017.12.005.
5. H. Bouma et al., "Automatic analysis of online image data for law enforcement agencies by concept detection and instance search," in *Proc. SPIE 10441, Counterterrorism, Crime Fighting, Forensics, and Surveillance Technologies*, 2017. doi: 10.1117/12.2277970.
6. Z. Lipton, "The mythos of model interpretability," *Commun. ACM*, vol. 61, no. 10, pp. 36–43, 2018.
7. S. Raaijmakers, M. Sappelli, and W. Kraaij, "Investigating the interpretability of hidden layers in deep text mining," in *Proc. 13th Int. Conf. Semantic Systems*, Amsterdam, The Netherlands, 2017, pp. 177–180.
8. N. T. Lee, "Detecting racial bias in algorithms and machine learning," *J. Inform., Commun. Ethics Soc.*, vol. 16, no. 3, pp. 252–260, 2018.
9. S. Sabour, N. Frosst, and G. Hinton, "Dynamic routing between capsules," in *Proc. 31st Conf. Neural Information Processing Systems (NIPS)*, 2017, pp. 3856–3866.
10. S. Corbett-Davies and S. Goel, The measure and mismeasure of fairness: A critical review of fair machine learning. 2018. [Online]. Available: https://arxiv.org/abs/1808.00023
11. R. Guidotti, A. Monreale, S. Ruggieri, F. Turini, F. Giannotti, and D. Pedreschi, "A survey of methods for explaining black box models," *ACM Comput. Surv.*, vol. 51, no. 5, 2018. doi: 10.1145/3236009.
12. F. C. Keil, "Explanation and understanding," *Annu. Rev. Psychol.*, vol. 57, no. 57, pp. 227–254, 2006.

Stephan Raaijmakers is a senior scientist at TNO, and, as of 2019, a professor of communicative artificial intelligence (AI) at Leiden University, The Netherlands, specializing in AI and natural language processing. He is currently the technical coordinator of two large European projects on security (ASGARD and TITANIUM). Contact him at stephan.raaijmakers@tno.nl.

This article originally appeared in
IEEE Security & Privacy, *vol. 17, no. 5, 2019.*

Robot Science Writers

Charles Day
The American Institute of Physics

I first learned of robot journalism about five years ago. A start-up called StatSheet had created software that took electronic versions of baseball scorecards and turned them into news reports. Known since 2011 as Automated Insights, the company and its competitors work with media outlets to automatically generate thousands of news stories a year.

For now, the stories are confined to game summaries, earnings reports, and other topics in which the principal content resides in readily harvestable data. But there are signs that robot writers are expanding their reach and becoming more capable. In a recent *Journal of Science Communication* editorial, science editor Mico Tatalovic discusses the prospects and implications of robot science writers (https://jcom.sissa.it/sites/default/files/documents/JCOM_1701_2018_E.pdf).

Writing about science might seem a tough challenge for a robot or, more accurately, an algorithm embodied in computer code. Science, especially physics, is rooted in abstract concepts and is replete with specialized terms. Among the new tools that Tatalovic describes is an artificial intelligence (AI) called Manuscript Writer, which the company sciNote released last November. The AI is a new feature in the company's electronic lab notebook, ELN. Once a scientist has all of his or her data and lab notes in ELN, Manuscript Writer steps in to write a report.

"The feature allows me to assemble and present data in a way that can lead to a publication with only minor modifications from me," is a quote from a happy user, Tessa Grabinski of the University of Michigan, that Tatalovic found on sciNote's website. The quote continues: "Not only does the new feature generate manuscripts quickly, it also provides several versions that can be used to assemble that perfect publication for your data."

Writing a scientific paper is not the same as writing a news story about it. Still, a new AI out of Columbia and Stanford universities called Science Surveyor can automatically establish a paper's context and significance, one of the key steps in science journalism. Fed an abstract and reference list, the AI trawls the corpus of scientific literature to return a list of experts, a plot of how the field has evolved, a reading list of background information, and an evaluation of whether the paper is in the mainstream or an outlier.

Why do Manuscript Writer and Science Surveyor succeed? The answer, I believe, is hinted at by Grabinski's quote. Ultimately, what usually counts in a scientific paper are results and theories embodied in plots and mathematical formulae.

ABOUT THE AUTHOR

Charles Day is *Physics Today*'s editor in chief. The views in this column are his own and not necessarily those of either *Physics Today* or its publisher, the American Institute of Physics.

This article originally appeared in
Computing in Science & Engineering, *vol. 20, no. 3, 2018.*

 Published by the IEEE Computer Society

Department

Empowering Extreme Automation via Zero-Touch Operations and GPU Parallelization

Jinan Fiaidhi and Sabah Mohammed
Lakehead University

Editor: Jinan Fiaidhi, Lakehead University;
jinan.fiaidhi@lakeheadu.ca.

The extreme automation model attracts increasingly more manufacturing enterprises to deploy their services and applications on the emerging automation infrastructure that come with extreme range of new requirements, including smart collaborative factories, personalized services with dramatic improvements in customer-experience, massive capacity, imperceptible latency, ultra-high reliability, global web-scale reach, and support for massive machine-to-machine communication. This automation infrastructure are being transformed into programmable, software-driven, service-based, and holistically managed infrastructures, utilizing enablers and catalysts, such as NFV, SDN, Edge Computing, and Internet of Every Thing. However, it is nontrivial to scale the manufacturing services automatically due to the dynamic nature of the extreme automation infrastructure and the dependencies among various manufacturing components. The ultimate challenge is to have an infrastructure with a scalable performance. Straight forward thinking may think of scalable performance in terms of adding additional processing capabilities to a manufacturing problem set or a simulation. This is where one would expect the problem to be solved faster by adding more CPUs and GPUs, but too often, when the other manufacturing constraints has not been considered, conversely, the performance will be the killing bottleneck. Because more parallelization means more communication and data movement between the independent services and tasks, the result often is even more communications between them. To embrace extreme automation we will need a new kind of collective communication that involve powerful management, analytics, and security capabilities to reduce the risk and mitigate the complexity—and manual labor—involved in deploying and managing multimanufacturing

Digital Object Identifier 10.1109/MITP.2019.2892162
Date of current version 27 March 2019.

environments and transitioning legacy architectures. The benefits of such collective communication include the following:

- *Cross-domain IT automation:* The ability of machines, devices, sensors, and people to connect and communicate with each other using tools like the extreme workflow composer (EWC) (https://www.extremenetworks.com/product/workflow-composer/) enables customized cross-domain, multimanufacturing IT automation, enabling full integration of compute, storage, security and network resources.
- *Decentralized Autonomous Decisions:* The ability of cyber-physical systems to make decisions on their own and perform tasks as autonomously as possible. Only in the case of exceptions, interferences, or conflicting goals are tasks delegated to a higher level. This means the ability to plug-and-play data center for manufacturing creation within seconds for a product of any size—with no need for additional software or servers. Customers can automate at their own pace with customizable workflows using popular tools such as Ansible (https://docs.ansible.com/ansible-tower/) and Rundeck (https://www.rundeck.com/ansible), as well as the ability to move to full IT automation with the use of tools like the EWC and the Extreme Management Center (https://www.extremenetworks.com/product/extreme-management-center/).
- *Information, Analytics and Data Transparency.* The ability of information systems to create a virtual copy of the physical world by enriching digital plant models with sensor data. This requires the aggregation of raw sensor data to higher value context information along with historical data to give better insights on whatever analytics required. The Tableau Embedded Analytics (https://www.tableau.com/embedded-analytics), for example, provides end-to-end application visibility and telemetry across data centers and VMs to enabling administrators supporting the business with real-time information to make informed decisions.
- *DevOps Integration:* DevOps transformed the software development process. It facilitates continuous delivery—that is to say faster and more efficient releases without a corresponding increase in operational risk. Devops is a methodology that was first introduced among the IT companies who developed Applications and Cloud services, with the aim to be more rapid, robust, and efficient in the launch of various software releases. This methodology is now becoming a priority for manufacturers as well, who are accompanying their machines with advanced control dashboards, predictive maintenance algorithms, and mobile applications, in order to monitor the machines themselves.[1]
- *Digital Cognitive Systems*. Cognitive systems transform data into smart data. It provides support to humans by aggregating and visualizing information to make informed decisions and solve urgent problems on short notice or the ability of cyber-physical systems to physically support humans by conducting a range of tasks that are unpleasant, too exhausting, or unsafe for their human coworkers. Cognitive systems processes are nonlinear, and they can process massive amounts of data, sometimes at a faster rate than the human brain. As big data accumulates, cognitive systems will be able to unlock the value buried in these massive data sets. Predictive and prescriptive analytics will be used to process data. An example of a vender advocating for developing cognitive systems for extreme automation is OpenText,[2] the Waterloo, Ontario, based startup that has grown to a position of global leadership in enterprise content management and is now providing cognitive enterprise platforms for humans to a creator of data platforms for humans and machines.

We are arguing in this column that all the above benefits cannot be achieved for a harmonized and effective extreme automation environment without the enforcement and the availability of following two notions:

1. *The Zero-Touch Provisioning (ZTP):* ZTP is the feature that allows the devices to be provisioned and configured automatically, eliminating most of the manual labor involved with a collective communication. ZTP allows the hardware or services to be installed directly

into the environment.[3] In the simplest form of ZTP, once the networking switch is powered on, it uses standard network protocols to fetch everything it needs for provisioning. It can send a DHCP query to get the proper IP address for connectivity and management, then use BootP/TFTP to get the right operating system image, and then use another TFTP request to get the right configuration file based on the application profile, which it downloads and runs. However, zero-touch based on the latest technology suggests more smart "hands-free" provisioning methods that are based on AI and ML that replaces the NO-OP Networking ZTP or the manual setup by technicians, as well as introducing new server-side remote setup. According to Gartner, these smart ZTP are called AIOps or Algorithmic IT Operations.[4] The new smart ZTP will require sophisticated autonomous systems and seamless interoperability to handle service requests and respond to communication events in real-time. The new AIOps platforms encompass the IT disciplines of performance management, service management, automation, and process improvement, along with technologies such as monitoring, service desk, capacity management, cloud computing, SaaS, mobility, IoT, and many more. The ZTP AIOps services cover three important provisioning aspects: environment visibility, predictive ability, and automation services.

GPUs are available everywhere in edge devices, for self-driving cars, on desktops and workstations, in data centers, in servers from all server builders (HPE, Dell, IBM, etc.), as well as in cloud services from every major cloud provider – including Google. A network trained in the cloud on a GPU can be deployed on any of these GPU-powered devices.

2. *Parallelization of GPUs based on the use general-purpose computing on graphics processing Units (GPGPU):* The graphics processing unit (GPU) has become an integral part of today's mainstream computing systems. Over the past six years, there has been a marked increase in the performance and capabilities of GPUs. The modern GPU is not only a powerful graphics engine but also a highly parallel programmable processor featuring peak arithmetic and memory bandwidth that substantially outpaces its CPU counterpart. The GPU's rapid increase in both programmability and capability has spawned a research community that has successfully mapped a broad range of computationally demanding, complex problems to the GPU. This effort in general-purpose computing on the GPU, also known as GPU computing, has positioned the GPU as a compelling alternative to traditional microprocessors in high-performance computer systems of the future.[5] These technology trends indicate that GPUs are departing from their traditional role as slave coprocessors and are becoming truly first-class computing elements, on par with CPUs. One can envision a system architecture where GPUs will run complete self-contained programs and will have full access to standard operating system services of their host system just like CPUs. The current progress in using GPUs is more than anyone can imagine with the acceleration reported by the various development workflows and applications across platforms including Caffe2, Cognitive

GPUs are typically programmed using a Single Program Multiple Data (SPMD) programming model, like NVIDIA, CUDA, or OpenCL. A SPMD model lets programmers spawn a large number of threads that execute the same program, although each thread can take a different control flow path. All these threads are organized into a computation grid of groups of thread blocks (i.e., groups of threads). Each thread block has an identifier, and each thread has an identifier within the thread block, which can be used by programmers to map the computation to the data structures.

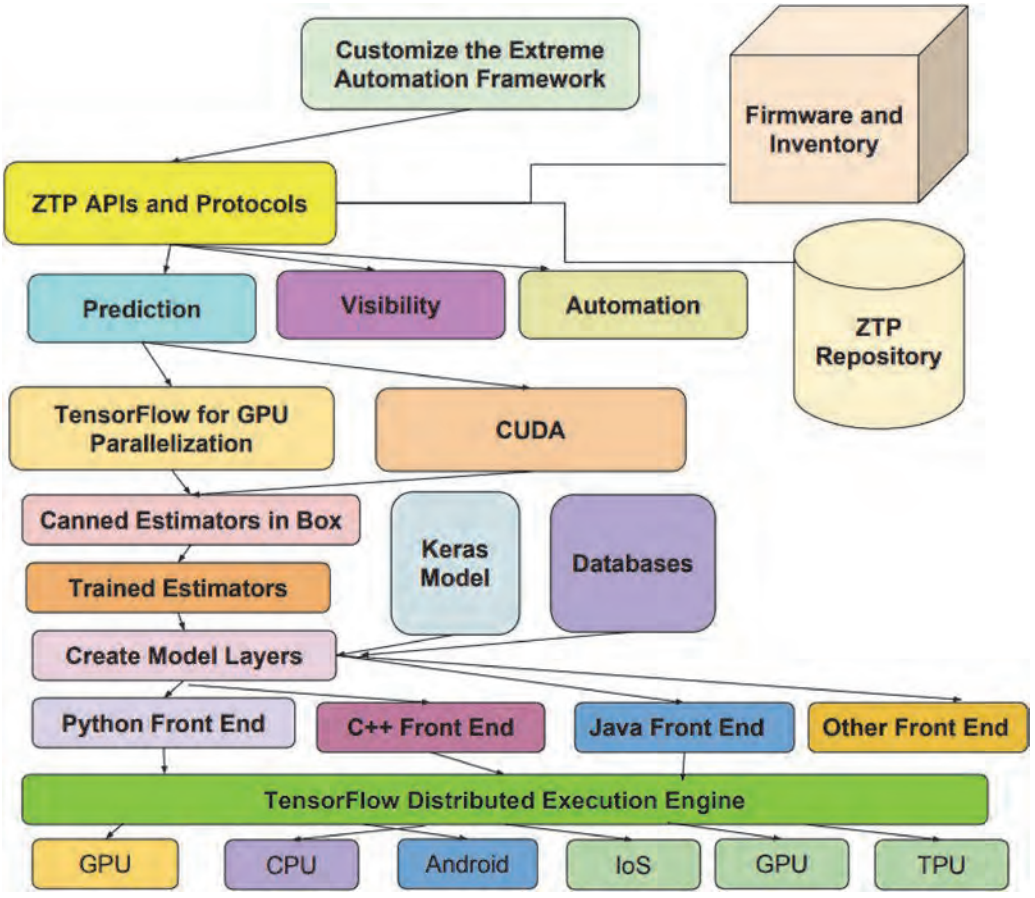

Figure 1. Predictive analytics for extreme automation utilizing ZTP and TensorFlow for systems with multiple GPUs.

Toolkit, Kaldi, MXNet, PaddlePaddle, Pytorch, and TensorFlow. Further performance improvement is also reported through what is generally known as GPU parallelization, which can improve performance not only by accelerating the targeted loops, but also by avoiding CPU-GPU communication. Achieving high applicability for parallelization on the GPU is critically important for performance because communication between CPU and GPU memories has high latency.[6] Many researchers have been developing a contemporary favorite for GPUs parallelism where they are utilizing GPGPU, a strategy exploiting the numerous processing cores found on high-end modern GPUs for the simultaneous execution of computationally expensive tasks. Modifying legacy CPU based algorithms to allow certain of their tasks to take advantage of GPU parallelization can demonstrate noteworthy gains in both task performance and completion speed.[7]

In[8], researchers were experimenting with ZTP like software called AMGE Modern, which is a provisioning networking interface and a compiler that support multi-GPU execution of computations like matrix multiplications. The AMGE imposes much lower memory footprint and coherence management overheads along with transparent data distribution that can be efficiently implemented on current GPUs using the UVAS and compiler/runtime-assisted code versioning. Using the array data type provided in AMGE results in shorter, faster, and cleaner code. AMGE achieves almost linear speedups for most of the benchmarks on a real 4-GPU system interconnected with moderate bandwidth. Further performance improvement can be achieved by reducing the virtual memory mapping granularity exposed by CUDA and by allowing programmers to tune the thread block scheduling policy. However, many other researchers showed similar performance improvements on real multiple GPUs for different algorithms like clustering.[9]

Actually the popularity of experimenting with GPUs is attributed to development environments like TensorFlow (https://www.tensorflow.org/), where it is becoming a popular platform for deep learning, which is quite useful in developing a predictive model for any extreme automation system. TensorFlow comes with lots of interesting features such as auto-differentiation (which saves you from estimating/coding the gradients of the cost functions) and GPU support. Alternatives to TensorFlow may include Torch, MXNet, Theano, Caffe, Deeplearning4j, and CNTK. Whatever the development environment that uses GPU, one can experiment with GPU parallelization that obviously requires ZTP provisioning protocol. In this direction, researchers are pointing to Keras API (https://keras.io/). Keras allows you to describe your GPU networks using high level concepts and write code that is backend agnostic, meaning that you can run the networks across different deep learning libraries. Typically, when you develop a multi-GPU training, you will need the Keras mutli_gpu_model that makes the parallel training/predictions easier (currently only available with TF backend). The main idea is that you pass your model through the method, and it is copied across different GPUs. The original input is split into chunks which are fed to the various GPUs and then they are aggregated as a single output. This method can be used for achieving parallel training and predictions. The combination of Keras and TensorFlow is currently used by the Google Cloud Platform[10] to perform machine learning images designed for deep learning practitioners. Programmers can use popular machine learning languages like Python or R to connect with Keras and Tensor Flow.[11] Figure 1 illustrates the way we may achieve Predictive Analytics, a common process for extreme automation, using the ZTP and the TenserFlow over systems that uses GPUs.

CONCLUSION

ZTP and deployment of GPU-based computing are the new terms for empowering extreme automation. With notions associated with these terms come new challenges and opportunities, as well as a dramatic shift in the landscape of programming the future collaborative manufacturing systems. We have only touched the surface of this exciting topic. We encourage you to contribute to this column by writing to the editor at jfiaidhi@lakeheadu.ca.

REFERENCES

1. J. Davis and R. Daniels, *Effective DevOps: Building a Culture of Collaboration, Affinity, and Tooling at Scale*. Sebastopol, CA, USA: O'Reilly, 2016.
2. M. Allen, OpenText in the machine era: Enterprise World 2017 Conf., InsightaaS Blog, Available: 2017-07-24, http://insightaas.com/opentext-in-the-machine-era/
3. E.H. Booth, III., W. M. Townsley, G. Weber, and W. Luo, "Techniques for zero touch provisioning of edge nodes for a virtual private network," U.S. Patent 7,535,856, issued May 19, 2009.
4. S. Shetty, *Gartner says Algorithmic IT Operations Drives Digital Business*, Gartner press, Release Apr. 11, 2017, [Online]. Available: https://www.gartner.com/en/newsroom/press-releases/2017-04-11-gartner-says-algorithmic-it-operations-drives-digital-business
5. J. D. Owens, M. Houston, D. Luebke, S. Green, J. E. Stone, and J. C. Phillips, "GPU computing," *Proc. IEEE*, vol. 96, no. 5, pp. 879–899, 2008.
6. I. Gelado, J. E. Stone, J. Cabezas, S. Patel, N. Navarro, and W.-M. W. Hwu, "An asymmetric distributed shared memory model for heterogeneous parallel systems," *ACM SIGARCH Comput. Archit. News*, vol. 38, no. 1, pp. 347–358, 2010.
7. G. Vasiliadis, M. Polychronakis, and S. Ioannidis, "Parallelization and characterization of pattern matching using GPUs," in *Proc. IEEE Int. Symp. Workload Characterization*, 2011, pp. 216–225.
8. J. Cabezas, L. Vilanova, I. Gelado, T. B. Jablin, N. Navarro, and W.-M. W. Hwu, "Automatic parallelization of kernels in shared-memory multi-gpu nodes," in *Proc. 29th ACM Int. Conf. Supercomput.*, 2015, pp. 3–13.
9. G. Andrade, G. Ramos, D. Madeira, R. Sachetto, R. Ferreira, and L. Rocha, "G-dbscan: A gpu accelerated algorithm for density-based clustering," *Procedia Comput. Sci.*, vol. 18, pp. 369–378, 2013.
10. V. Kovalevskyi, Deep Learning Images For Google Compute Engine, The Definitive Guide, Deep Learing As I See IT Blog, Jul. 3, 2018, [Online]. Available: https://blog.kovalevskyi.com/deep-learning-images-for-google-cloud-engine-the-definitive-guide-bc74f5fb02bc
11. R Views, "Connecting R to Keras and TensorFlow," Dec. 10, 2017, the R Bloggers. [Online]. Available: https://www.r-bloggers.com/connecting-r-to-keras-and-tensorflow/

Jinan Fiaidhi is a Full Professor of computer science and the Graduate Coordinator of the BioTech Ph.D. program with Lakehead University, Thunder Bay, ON, Canada. She is also an Adjunct Research Professor with the University of Western Ontario, London, ON, Canada, and the Editor in Chief of the IGI Global *International Journal of Extreme Automation and Connectivity in Healthcare*. She is also the Chair of Big Data for eHealth with the IEEE ComSoc. Contact her at jfiaidhi@lakeheadu.ca.

Sabah Mohammed is a Full Professor with the Department of Computer Science and Supervisor with the Smart Health FabLab, Lakehead University, Thunder Bay, ON, Canada. He is also an Adjunct Professor with the University of Western Ontario, London, ON, Canada. Moreover, he is the Chair of Smart and Connected Health with the IEEE ComSoc. Contact him at mohammed@lakeheadu.ca.

This article originally appeared in
IT Professional, *vol. 21, no. 2, 2019.*

COMPSAC 2020
Madrid, Spain
July 13-17, 2020

COMPSAC is the IEEE Computer Society Signature Conference on Computers, Software and Applications. It is a major international forum for academia, industry, and government to discuss research results and advancements, emerging challenges, and future trends in computer and software technologies and applications. The theme of COMPSAC 2020 is "Driving Intelligent Transformation of the Digital World".

Staying relevant in a constantly evolving digital landscape is a challenge faced by researchers, developers, and producers in virtually every industry and area of study. Once limited to software-enabled devices, the ubiquity of digitally-enabled systems makes this challenge a universal issue. Furthermore, as relevance fuels change, many influencers will offer solutions that benefit their own priorities. Fortunately, history has shown that the building blocks of digital change are forged by those conducting foundational research and development of digital systems and human interactions. Artificial Intelligence is not new, but is much more utilized in everyday computing now that data and processing resources are more economically viable, hence widely available. The opportunity to drive the use of this powerful tool in transforming the digital world is yours. Will your results help define the path ahead, or will you relegate those decisions to those with different priorities for utilizing intelligence in digital systems? COMPSAC has been and continues to be a highly respected venue for the dissemination of key research on computer and software systems and applications, and has influenced fundamental developments in these fields for over 40 years. COMPSAC 2020 is your opportunity to add your mark to this ongoing journey, and we highly encourage your submission!

COMPSAC 2020, organized as a tightly integrated union of symposia, will focus on technical aspects of issues relevant to intelligent transformation of the digital world. The technical program will include keynote addresses, research papers, industrial case studies, fast abstracts, a doctoral symposium, poster sessions, and workshops and tutorials on emerging and important topics related to the conference theme. Highlights of the conference will include plenary and specialized panels that will address the technical challenges facing researchers and practitioners who are driving fundamental changes in intelligent systems and applications. Panels will also address cultural and societal challenges for a society whose members must continue to learn to live, work, and play in the environments the technologies produce.

Authors are invited to submit original, unpublished research work, as well as industrial practice reports. Simultaneous submission to other publication venues is not permitted except as highlighted in the COMPSAC 2020 J1C2 & C1J2 program. All submissions must adhere to IEEE Publishing Policies, and will be vetted through the IEEE CrossCheck portal. Further info is available at www.compsac.org.

Organizers
Standing Committee Chair: Sorel Reisman (California State University, USA)
Steering Committee Chair: Sheikh Iqbal Ahamed (Marquette University, USA)
General Chairs: Mohammad Zulkernine (Queen's University, Canada), Edmundo Tovar (Universidad Politécnica de Madrid, Spain), Hironori Kasahara (Waseda University, Japan)
Program Chairs in Chief: W. K. Chan (City University, Hong Kong), Bill Claycomb (Carnegie Mellon University, USA), Hiroki Takakura (National Institute of Informatics, Japan)
Workshop Chairs: Ji-Jiang Yang (Tsinghua University, USA), Yuuichi Teranishi (National Institute of Information and Communications Technology, Japan), Dave Towey (University of Nottingham Ningbo China, China), Sergio Segura (University of Seville, Spain)
Local Chairs: Sergio Martin (UNED, Spain), Manuel Castro (UNED, Spain)

Important Dates
Workshops proposals due: 15 November 2019
Workshops acceptance notification: 15 December 2019
Main conference papers due: 20 January 2020
Paper notification: 3 April 2020
Workshop papers due: 9 April 2020
Workshop paper notifications: 1 May 2020
Camera-ready and registration due: 15 May 2020

IEEE COMPUTER SOCIETY

Photo: King Felipe III in Major Square, Madrid
Photo credit: Iria Castro - Photographer (Instagram @iriacastrophoto)

Department: Extreme Automation
Editor: Jinan Fiaidhi, jinan.fiaidhi@lakeheadu.ca

Security and Vulnerability of Extreme Automation Systems: The IoMT and IoA Case Studies

Jinan Fiaidhi and Sabah Mohammed
Lakehead University

One of the key advantages of extreme automation is its digital transformation speed to multiple businesses, including manufacturing, healthcare, and aviation, which allows for rapid communication, iteration, and sharing of services and their corresponding physical representation. While this enables a more efficient automation process, it also presents opportunities for cyber-attacks and mitigated risks by impacting the physical word. According to Frost and Sullivan forecast, the growth in the Internet of Medical Things (IoMT) will reach $72.02 billion by 2021 with more than 30 billion connected medical devices in the healthcare ecosystem (http://ww2.frost.com). Nearly three in five healthcare companies now have important IoMT medical devices installed, with seven out of eight strategizing its use.[1] GE Healthcare, Medtech, Medtronic, and Philips, and even technology giants such as Apple, IBM, Cisco, and Qualcomm, are developing capabilities in IoMT applications with a variety of connected medical devices from pregnancy testing kits to Sugar IQ diabetes assistant. Actually, about 70% of healthcare organizations now utilize the technology for maintenance and monitoring practices. Nearly 64% of the technology's use is dedicated to patient monitoring. The technology is allowing these devices to generate, collect, analyze and transmit data, creating the IoMT—a connected infrastructure of health systems and services. IoMT devices form a connected ecosystem of sensors and devices tagged around the patient to capture, measure, and identify key data; stratify risks; make decisions; and initiate necessary action plans. The healthcare industry is on the verge of becoming more proactive patient care by leveraging and embracing the IoMT initiative. We are seeing impressive connected medical devices and

Digital Object Identifier 10.1109/MITP.2019.2906442
Date of current version 17 July 2019.

 Published by the IEEE Computer Society

applications in the healthcare setting such as the following:

- the deployment of hospital RFID, beacon or indoor GPS technologies for navigating within the premises;
- virtual assistants for homecare to help patients and seniors with their self-care using mHealth applications and smart diagnostic medical devices that support telehealth services;
- smart vehicles that can track vitals of passengers during transit;
- exigency support by drones for emergency response;
- smart, digitized clinical devices like digital stethoscopes for clinicians in primary care;
- smart hospital rooms that allow patients to communicate with care teams virtually from their bedside;
- kiosks at community centers to improve access to informational services, pharmaceutical products, and telemedicine services.

Beyond the concepts explained above, the IoMT is fascinating because its relative infancy means that many regulatory and compliance controls and practices are still in their nascent stages. Highly influential agencies in the healthcare industry, such as the FDA, are understandably still determining how to help guide organizations to ensure that those organizations—and the patients they serve—are getting the most out of the IoMT.

A similar technological shift is happening in the aviation front. The doubling of air traffic by 2020 and high CO_2 emissions are among the challenges that the aviation industry must address to become financially and environmentally sustainable. Connectivity and the Internet of Things provide a cost-effective opportunity to tackle these challenges, enabling an improved passenger experience while increasing operational efficiency and safety. The Internet of Aviation (IoA), or Aviation 4.0, is the door for solving these challenges in the next future to unlock the potential of IoT and leverage it for competitive advantage. For airlines, the IoA offers a unique opportunity to deliver new value to travelers, cargo customers, and shareholders. Airlines will have a constant, extremely detailed, live picture of their network. This will include the aircraft functionality, the crew activities, the gate situations, the airline hub, and their schedule information, all employed for more effective airline management. Just think about the aircraft itself, which has more than 5000 sensors and generates up to 10 GB of data per second,[2] as the availability of this data will enable a truly connected airline to be a smart airline as well, deliver exceptional customer services, and win in the marketplace. With extremely connected aviation devices entering the picture, the nature and the rigor of the aviation work changes. As IoA assumes more control of the craft, the pilot is relieved of much manual labor. This evolution comes with a rise of information to be managed by the pilot, who might be confronted with more than 600 devices and indicators to be monitored and controlled in the cockpit. Actually, the airline industry company Honeywell took its Boeing 757 IoT-connected test aircraft on a tour around the world in May 2017[3] with greater promise for an excellent flying experience for pilots, passengers, and aircraft operators. The civil aviation authorities are recommending the use of systems like the aircraft health monitoring system (AHMS) to reduce the possibilities of having any fault on flight systems. Systems like the AHMS bring vast improvements to the utilization and analysis of the aircraft data to enhance availability, reliability, and safety of the aircraft, which in turn drives the take up of condition-based maintenance projects to streamline maintenance, repair, and overhaul. However, when more advanced systems begin to take over planning and analysis functions, such as setting and adjusting a flight plan, the pilot becomes less engaged not only physically but mentally.[4] The recent fatal crushes of Ethiopian Airlines Flight 302 and Lion Air Flight 610 shows similar patterns of faults in the automatic control of these flights.

Because of increasing demand and its accessibility to high internet speed, IoMT, and IoA has opened doors for serious vulnerabilities to healthcare and aviation systems. The disastrous consequences of these issues will not only disrupt services causing financial losses but will also put the peoples' lives at risk.

IOMT ADOPTION VULNERABILITIES

We need to understand that what makes our devices smart is that they are embedded with

little computers or controllers that are connected to the Internet. It is this connectivity that enables the flow of communication between the device and a backend that collects the data and helps the user or other devices to do something more efficiently. These devices are essentially the physical manifestation of applications like we have on our phones or the web. From a security perspective, hackers can attack these devices much in the same way that they would exploit a traditional endpoint device like a desktop computer. This means that with the right vulnerability, they can take it over, access its data, overwhelm it, or perform other kinds of malicious activities. However, unlike your desktop computer, there are some key differences that need to be considered. However, the accelerated adoption of IoMT requires careful considerations as they must be robust against the security vulnerabilities affecting medical devices that IoMT uses, a landscape of uncertain liability, new standards, and emerging policies and regulations. Consequently, medical device manufacturers should keep abreast of current minimum security standards to prevent cyberattacks like the "WannaCry" ransomware attack in May 2017.[5] These security vulnerabilities highlight the importance of developing standards, using best practices for compliance. The existing broad, ambiguous standards regulating the IoMT invite litigation, and precise legal boundaries have yet to be drawn. In an effort to regulate the IoMT and ensure public safety, the US Food and Drug Administration (FDA) issued premarket and postmarket cybersecurity guidance in October 2018, providing nonbinding recommendations to device manufacturers.[6] However, many question remains to be answered regarding IoMT standards and security such as[7] the following:

- What is the reasonable standard of care in creating a secure IoMT device?
- What constitutes a design defect or failure to warn?
- Are security vulnerabilities considered a design defect?
- For how long must device manufacturers provide security monitoring and software updates after selling a product?
- Does user failure to download security updates act as a superseding cause or a failure to mitigate in cases of liability for defective software?
- Will these security vulnerabilities mean an uptick in shareholder derivative actions?

Craig Badrick estimates that of every 1000 IoT devices in use, 164 are subject to attacks. As hospitals discover increasingly more applications for the IoMT, many are beginning to add devices that may actually be putting their operations—and even patient lives—at risk.[8] In fact, the healthcare industry is rapidly moving to a completely digitized environment, and as a result, devices have been introduced to the hospital ecosystem and bedside workflows to help extend and streamline care throughout the hospital, and many devices are incorporated to monitor remotely patients at home or work while using robust medical devices and mobile smartphones have allowed clinicians to become more efficient and mobile with patient care. Unfortunately, this new technology has also opened the door to increased risk and new potential points of exposure for healthcare IT infrastructures. Without enforcing rigorous standards for safely using these medical devices within the new IoMT platform, then each network-connected medical device within a health provider's ecosystem will open up the possibility for patient health information exposure, as well as the potential for other unauthorized use of critical systems and applications.

Other vulnerabilities include interoperability between medical devices, regulatory changes, and scalability.[9] The cybersecurity challenge, however, is the most severe one of all as healthcare is promoting more mobile medical devices in their setting. Every medical device then becomes like a "back door" into a hospital's IT network, and attackers are now exploring a new strategy called "destruction of service," or DeOS, which will completely incapacitate the network. This issue needs further consideration as the healthcare setting includes a variety of easily compromised medical devices—some of the medical devices were built 20 years ago but still work, believe it or not. Sadly, many of these tools are still being used by hospitals (often to save money). However, those medical devices, including pacemakers, X-ray machines, and CT scanners, use outdated security software that isn't automatically updated. This leaves hospitals and patients very vulnerable. Moreover, regulatory agencies like the FDA-issued

security recommendations are not mandates—quite simply, without a firm mandate to follow, manufacturers and healthcare organizations struggling to follow the FDA's did not provide decisive guidelines to reduce device security risks, especially if it costs more money and resources. According to a 2017 study conducted by the Ponemon Institute, only 51% of medical device manufacturers and 44% of healthcare organizations follow FDA guidelines.[10] The pacemaker recall of 465 000 devices by the FDA in August 2017 is an example of the cybersecurity vulnerabilities on these sensitive medical devices where it could allow a hacker to take over the medical device that controls heart rhythm.[11] Actually, the IoMT has extended the traditional perimeter of healthcare security. Once a threat is successfully inside, there are usually few security measures in place to detect it or slow it down. This is one reason why IoMT devices are popular attack vectors. These internal endpoints have been authorized to access the network as an authorized user. Once deployed inside the perimeter defenses, IoMT devices have largely unquestioned access to the network's data.[12] These vulnerabilities will never be fully prevented as the technology never stays still, and neither do hackers. The new developments in protecting patient data and patients themselves will not be the end of healthcare cybersecurity, nor will they guard against every possible way of hacking IoMT devices. However, it is important to develop and implement medical device security and IoMT security strategies. These strategies need to include not only a screening and threat mitigation standard for current devices but also a plan for maintaining security on a continuing basis. These strategies must not disrupt clinical workflow, but it should provide clinicians with the proper knowledge on what to do if a data compromise occurs. Figure 1 illustrates the types of security attacks that may be imposed on IoMT devices and networks. The outage attack stops an IoMT device (e.g., pacemaker) from working and may result in death or physical injury to the patient. With physical attacks, hackers need to physically install a pseudo sensor in the IoMT architecture in order to receive unauthorized health information. The hacker may physically alter the device to state false readings, resulting in the patient receiving heavier medication, for example. The message

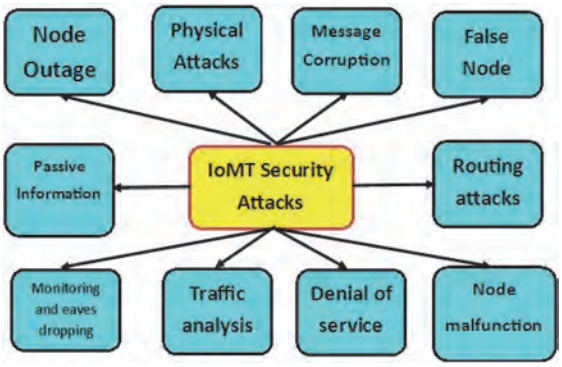

Figure 1. Types of Security Attacks on IoMT.

corruption attack is the result of inserting a virus with IoMT data when sent to the physician causing corruption of the original data. In the false node attack, one patient data may be replaced with another patient data and would unknowingly result in an inaccurate diagnosis to both patients. In a passive information gathering attack, the hacker would collect the data as they are sent to clinicians as an intermediary and may store all patient data in a geographical area and sell that data to insurance companies or other beneficiaries. In a routing attack, the hacker could create an infinite loop between the various sensors in the IoMT network, and the data would constantly overwrite itself. In a monitoring and eavesdropping attack, the hacker could use the wearable IoMT device to track the patient's voice commands and listen to personal conversations. Those conversations could be recorded and used to blackmail the patient. A traffic analysis attack happens when a patient sends data from their IoMT device to their family or friends, where in the hacker could send additional messages to the recipients. In a Denial of service attack, the hacker may lock the medical device with a password encryption and prevent a patient from using it. The hacker would provide the password only if the ransom is paid. Finally, in a node malfunction attack, the hacker may erase emergency phone numbers on an emergency response sensor and prevent patients from calling emergency services.

There are many popular brands that be subjects to these attacks, especially if they were used within the healthcare IoMT settings. Examples may include Medtronic MyCareLink, which is an app that uses a reader to receive pacemaker data. The data are sent from the Smartphone to the

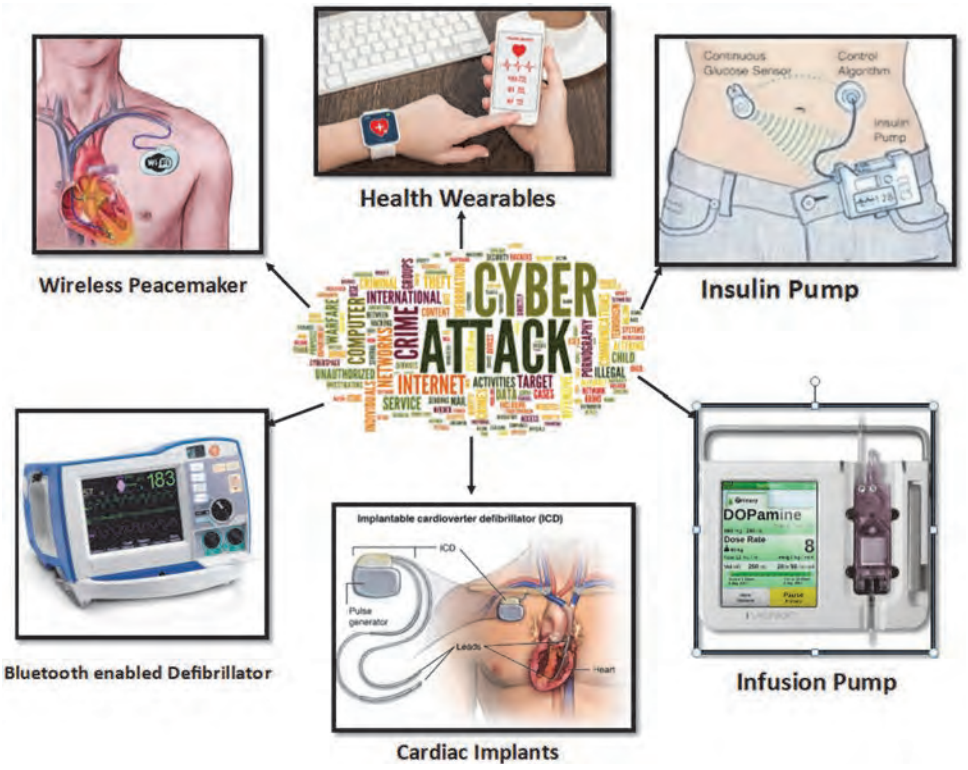

Figure 2. Type of IoMT devices most vulnerable to security attacks.

patient's clinic. The Reveal LINQ is another example which uses a tiny insertable monitor placed just under the skin and the MyCareLink Patient Monitor—a bedside unit that collects heart rhythm data from the insert and sends it to your doctor. The OmniPod Insulin Management System is a third example which includes a tubeless Pod and a handheld Personal Diabetes Manager that allows the patient to wireless program insulin delivery. The CardioNet wEvent is a fourth example that uses a wireless cardiac event monitor to collect asymptomatic and symptomatic events to detect heart arrhythmias. These data are automatically transmitted wirelessly through a cellular network to the physician. Finally, the Welch Allyn Home Blood Pressure Monitor with the Blood Pressure App sends readings directly to doctors and stores them to track results. Figure 2 illustrates the most vulnerable IoMT devices, as investigated in reference.[13]

DEFENDING IOMT FROM SECURITY ATTACKS

Hacker meddling in the IoMT operations not only costs lots of time, money, and operational downtime, but threatens lives. It is a dire situation that must be addressed. Hospitals and other healthcare providers must practice better cybersecurity hygiene. For starters, healthcare organizations must improve the speed and thoroughness of software patching and update processes. As much as possible, organizations also need to use threat intelligence and automation, as well as institute cyber-awareness training programs to protect against social media attacks and other attack vectors. According to Adefala,[14] there are some foundational activities that healthcare organizations can take to ensure they are protected against such attacks:

- Practice diligent cyber hygiene.
- Reinforce network segmentation.
- Achieve transparent visibility and control.
- Use advanced threat intelligence.

IOA VULNERABILITIES

While flying has always been one of the safest ways to travel, thanks to its wide-ranging international regulatory frameworks, aviation incidents have an outsize impact on the public

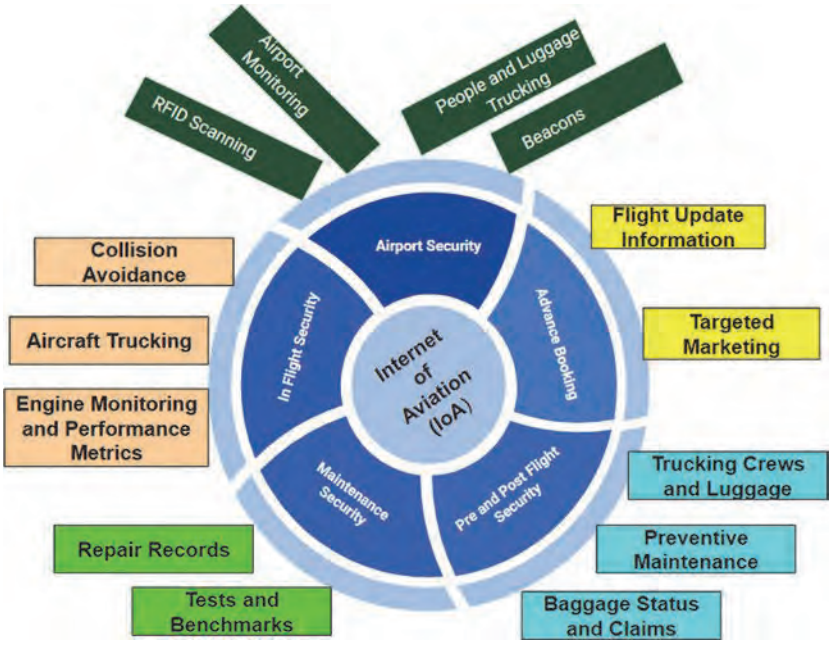

Figure 3. Major security sections for the IoA technology.

consciousness. From recent airport attacks to the crash of several flights like the Ethiopian Airlines Flight 302, horrifying images are more powerful than reassuring statistics. Figure 3 illustrates the major types of security involved with IoA.

IoA technologies (automation, IOT, artificial intelligence, cognitive computing, streaming analytics, digitization, etc.) have the potential to generate a paradigm shift in the aviation industry, generating new mechanisms to make it not only more efficient but also safer. Unexplored concepts and approaches to aviation vulnerability and safety need to start on different fronts like[2]

- automatic flying in predefined situations in a rule-based way;
- developing a robust aircraft predictive maintenance;
- cockpit safety cognitive computing aid systems;
- real-time weather information update;
- improved search and rescue services especially in the oceanic or remote area;
- real-time human performance monitoring and alerting based on nonintrusive physiological sensors/signals and contextual information;
- worldwide aeronautical networks interoperability, including signal processing and wireless performance, as well as the aircraft interfaces to the Internet;
- verification and validation of the onboard software, how to secure end-to-end entire SW supply processes, and the understanding of cyber-physical life-cycle scale;
- improvement of airplane health, control, and prognostics by exploiting sensor networks and data fusion, information management and data analytics and, critical real-time data sharing, appropriate end-to-end information exchange, distributed decision-making;
- human-automation interface issues such as visualization, keeping human-in-the-loop and connection between aircraft controls and air traffic systems.

According to the recent IATA report, the next 30 years are likely to be more turbulent, as a new wave of technological change and innovation unfurls. Some see this wave sweeping the airline industry, citing as precedents the taxi industry before Uber arrived, the music industry before internet downloads, and the printing industry before computer design software.[15] Disruption to the global transportation network, for example, can cause ripples of economic and social turmoil. It is little wonder that cybersecurity is ranked as the number one challenge in the air transport industry.

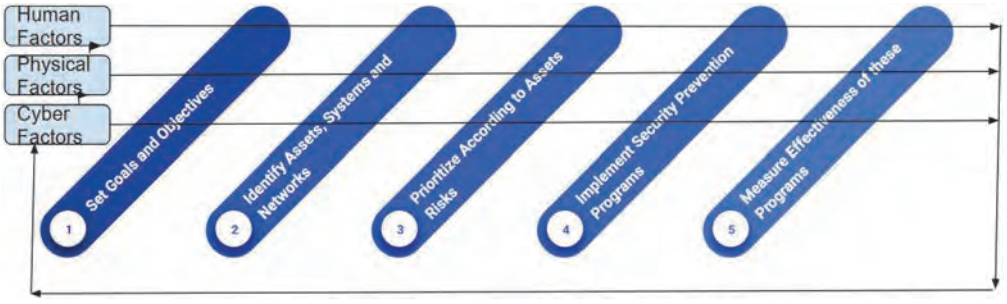

Figure 4. Recommended boeing iterative security prevention cycle.

DEFENDING IOA FROM SECURITY ATTACKS

According to Boeing preventing security attacks in its core is an iterative process that every aviation industry must practice.[16] Figure 4 illustrates this security prevention iterative strategy.

Although Boeing has a lot of experience with preventing security factors, the problem remains in the way we develop our software for mission-critical systems such as IoMT and IoA. Lamport[17] provided a magic solution for all the problems associated with software development using a metaphor that he calls *why we should build software like we build houses*. Programs need to be designed like the construction blueprints where everything is transparent and robust. This issue is the root of vulnerability of systems with extreme connectivity and collaboration. We leave this issue to Academia and the IoMT and IoT industries and let us use these blueprint languages in programming these connected devices and services.

IoMT and IoA pass though a massive wave of digitization change in order to make both industries affordable, safe, and smart. This article has shed the light on the security and vulnerability issues of these two technologies and has suggested such remedies as echoed by the relevant industries.

CONCLUSION

IoMT and IoA are emerging waves of technologies that contributes to establishing-connected systems. It consists of smart devices, such as wearables, sensors technology, smart algorithms, and monitors, strictly for healthcare and aviation uses. It can reduce unnecessary hospital visits and the burden on aviation systems. However, we have only touched the surface of this exciting topic. We encourage you to contribute to this column by writing to the editor at jfiaidhi@lakeheadu.ca.

REFERENCES

1. K. Matthews, "Is the internet of medical things (IoMT) on par with the IoT market as a whole?," Hit Consultant, 01/30/2019. [Online]. Available: https://hitconsultant.net/2019/01/30/is-iomt-tech-iot-market-as-a-whole/#.XI0plihKiUk
2. R. A. Valdés, V. F. G. Comendador, A. R. Sanz, and J. P. Castán, "Aviation 4.0: More safety through automation and digitization, Aircraft Technol., Melih Cemal Kuşhan, IntechOpen, Mar. 9, 2018, doi: 10.5772/intechopen.73688. Available: https://www.intechopen.com/books/aircraft-technology/aviation-4-0-more-safety-through-automation-and-digitization
3. F. Roberts, "Honeywell's IoT connected aircraft takes flight, Internet of business," Jun. 19, 2017. [Online]. Available: https://internetofbusiness.com/honeywells-iot-aircraft-flight/
4. N. Carr, On autopilot: The dangers of overautomation, ROUGH TYPE, Mar. 14, 2019. [Online]. Available: http://www.roughtype.com/?p=8622
5. Z. Rodionova, "Healthcare is now top industry for cyberattacks, "*Independent*," Apr. 21, 2016. [Online]. Available: https://www.independent.co.uk/news/business/news/healthcare-is-now-top-industry-for-cyberattacks-says-ibm-a6994526.html
6. S. Schwartz, "Premarket submissions for management of cybersecurity in medical devices, draft guidance for industry and staff, food and drug administration (FDA)", Oct. 18, 2018. [Online]. Available: https://www.fda.gov/downloads/MedicalDevices/DeviceRegulationandGuidance/GuidanceDocuments/UCM623529.pdf

7. M. Segura, C. M. Butler, F. Tabibkhoei, and R. Smith, "The internet of medical things raises novel compliance challenges, medical devices online," Jan. 3, 2018. [Online]. Available: https://www.meddeviceonline.com/doc/the-internet-of-medical-things-raises-novel-compliance-challenges-0001
8. C. Badrick, "Best practices in IoMT security, turn key technologies," Jan. 4, 2019. [Online]. Available: http://www.turn-keytechnologies.com/blog/network-solutions/best-practices-iomt-security
9. J. Haughey, K. Taylor, M. Dohrmann, and G. Snyder, "Medtech and the internet of medical things: How connected medical devices are transforming health care," Deloitte 2018. [Online]. Available: https://www2.deloitte.com/global/en/pages/life-sciences-and-healthcare/articles/medtech-internet-of-medical-things.html
10. Ponemon Instit., "Medical device security: An industry under attack and unprepared to defend," White Paper, May 2017. [Online]. Available: https://www.synopsys.com/content/dam/synopsys/sig-assets/reports/medical-device-security-ponemon-synopsys.pdf
11. A. Young, "What internet of medical things (IoMT) devices mean for healthcare cybersecurity vulnerable IoMT devices," Apr. 2018. [Online]. Available: https://healthtechmagazine.net/article/2018/04/The-Internet-of-Medical-Things-Opens-Health-Organizations-Up-to-More-Threats
12. J. Nguyen-Duy, "Healthcare's secret weapon for securing the IoMT, CSO report," Feb. 1, 2018. [Online]. Available: https://www.csoonline.com/article/3252150/healthcare-s-secret-weapon-for-securing-the-iomt.html
13. M. Anandarajan and S. Malik, "Protecting the internet of medical things: A situational crime-prevention approach, cogent medicine," *Cogent Med.*, vol. 5, 2018, Art. no. 1513349. [Online]. Available: https://www.cogentoa.com/article/10.1080/2331205X.2018.1513349.pdf
14. L. Adefala, "Healthcare experiences twice the number of cyber attacks as other industries," CSO ONLINE, Mar. 6, 2018. [Online]. Available: https://www.csoonline.com/article/3260191/healthcare-experiences-twice-the-number-of-cyber-attacks-as-other-industries.html
15. P. Steele, "Future of the airline industry, 2035," IATA 2018. [Online]. Available: https://www.iata.org/policy/Documents/iata-future-airline-industry-pdf.pdf
16. R. Rencher, "Boeing, Securing arline information on the ground and in the air," AERO QTR_03.12, 2012. [Online]. Available: https://www.boeing.com/commercial/aeromagazine/articles/2012_q3/5/
17. L. Lamport, "Why we should build software like we build houses," Jan. 2013. [Online]. Available: https://www.wired.com/2013/01/code-bugs-programming-why-we-need-specs/

Jinan Fiaidhi is a Full Professor of computer science and the Graduate Coordinator of the BioTech Ph.D. program with Lakehead University, Thunder Bay, ON, Canada. She is also an Adjunct Research Professor with the University of Western Ontario and the Editor-in-Chief of the IGI Global *International Journal of Extreme Automation and Connectivity in Healthcare*. She is also the Chair of Big Data for eHealth with the IEEE Communications Society. Contact her at jfiaidhi@lakeheadu.ca.

Sabah Mohammed is a Full Professor with the Department of Computer Science and Supervisor of the Smart Health FabLab, Lakehead University, Thunder Bay, ON, Canada. He is also an Adjunct Professor with the University of Western Ontario, London, ON, Canada. Moreover, he is the Chair of Smart and Connected Health with the IEEE Communications Society. Contact him at mohammed@lakeheadu.ca.

This article originally appeared in
IT Professional, *vol. 21, no. 4, 2019.*

Going Back to the Roots—the Evolution of Edge Computing, an IoT Perspective

Marjan Gusev
Sts. Cyril and Methodius University

Schahram Dustdar
TU Wien Distributed Systems Group

When the requirements and processing appetites grew sufficiently in the era of personal computing, users began to think about cloud servers. This initiated the growth of cloud computing and set up the era of Everything-as-a-Service. When the IoT idea spread, lots of sensors and smaller devices were introduced for ubiquitous and pervasive computing. The simplest way to collect and process all this data was to connect the devices to the cloud. This required high-speed networks and interconnections. However, no matter how fast they could go, engineers understood that there's no need to transfer all the data through links to the central data server, and that it would be much more efficient to distribute smaller servers in front of the central data server and closer to the location of the user. This began the evolution of pushing data collection and processing back to end-user devices, as a new computing approach called edge computing, which appears in different forms such as fog computing, cloudlets, and mobile edge computing.

The Internet of Things (IoT) is a network of various devices that communicate with each other via the Internet. The proliferation of these devices in a communication-actuating network creates the IoT, wherein sensors and actuators blend seamlessly with the environment around us, and the information is shared across platforms.[1]

Various hardware, software, data, and services are interconnected and the application domain is huge, including smart environments (smart homes, smart cities, etc.), healthcare (triage, patient monitoring, health status, etc.), emergency services (remote monitoring, resource management,

 Published by the IEEE Computer Society

etc.), transport control (traffic management, infrastructure monitoring, etc.), and environment-related services (resource supply, environment monitoring, etc.).

The interconnected ”things” have various characteristics according to their location, type, nature, power supply, and various other aspects, including, for example, characteristics of associated data that represent their state. Analyzing the location, they can be static or movable devices that need a mobile and wireless connection to enable freedom of movement—for example, wearable sensors are those that are worn on a human body and are battery operated using mobile and wireless to transmit data.

The location of a “thing” also dictates the power resource, whether it be a standard electrical power supply or battery. Battery-operated devices require small processing capabilities due to limited battery life, which raise the need to offload storage and processing to a more powerful server, such as cloud and fog computing. The data collected and processed is huge, which raises the need for offloading to a server. The offloading idea entails new architecture issues, such as where the server should be located, how to avoid long communication delays, and how to ensure low energy consumption and enable greater mobility and independence of IoT devices. This paper will analyze the evolution of communication, storage, processing, and energy consumption requirements in the IoT world. The requirements analysis is elaborated with appropriate architectural designs of computer infrastructure.

BACKGROUND AND BASIC DEFINITIONS

Sensors and actuators are essential IoT devices for production (collection) of data or control devices. Our goal is to realize small IoT devices independent of the platform, location, and environment. This can be achieved by offloading the storage and computation from the IoT devices to servers or edge devices. It will allow the IoT devices to be mobile and wirelessly connected to the environment.

In this context, edge refers to end-user devices located close to the IoT devices. Servers that communicate to edge devices, often called edge servers, are those set in proximity of the edge devices, which avoids communication latency. Examples include smaller servers at base stations or sometimes, mobile phones.

Clouds are one way to realize the offloading idea, where edge computing is an extension to clouds closer to the end-devices. Edge computing refers to the enabling technologies that allow computation to be performed at the edge of the network on downstream data on behalf of cloud services and upstream data on behalf of IoT services.[2]

The emergence of fog computing aims to take advantage of positioning servers close to base stations and cope with increased data traffic, and therefore, hosting services, workloads, applications, and large amounts of data at the edge of the network.[3,4] This makes fog computing a synonym for edge computing because it facilitates the operation of computing, storage and networking services between end devices and cloud computing data centers. This approach is the communication community’s answer to the rise of IT business to deliver IT-based services.

One well-known challenge of using the cloud as a server is the long latency between the mobile device and the cloud server in comparison to localized computing and small-scale distributed computing called cloudlets.[5] In collaboration with the cloud, cloudlets make recognition and afterward send results back to the mobile device. In a cloudlet solution, mobile devices directly exchange data with the local server on the cloudlet using WiFi, instead of connecting to the cloud server.[6]

Mobile edge computing is the realization of edge computing at the base stations of mobile operators.[7] Thus, it is a special kind of edge computing with precisely located intermediate servers.[8] Cloudlets are also smaller servers between end-user devices and servers located on the premises of Internet providers in their LAN.

ICT PROCESSING REQUIREMENTS AND SOLUTIONS FOR IOT

Modern technology trends and available technology are pushing the commercialization and mass exploitation of various IoT sensors and remote controller devices.

Basic requirements

Let's begin the analysis with a description of a typical IoT sensor or remote controller device. An IoT sensor is a device capable of sensing data such as environmental information or digital data from social networks. It detects changes from any material (physical, chemical, or biological) and some non-material information. The sensor includes the following functions: sensing, transforming information into an electrical signal, and transmitting the electrical signal to the neighboring environment.

Some IoT sensors are connected to an IoT controller device where sufficient information that will trigger an action on the controller device. It can be any other IoT device that is not sensing and is responsible for controlling (moving) a mechanism or system, usually called an actuator.

The communication between the IoT sensor and controller device may be realized by a personal area network connection, usually realized via a direct cable, Bluetooth, or other wireless connection. This forms an independent autonomous sensing and controller system. Unless this system memorizes the previous states, there is only a limited need for processing using small amounts of memory storage.

Requirements to offload computations and storage

During the time, the requirements and processing appetites have increased and engineers have tried to realize more complex systems and connect them to the outside world. The increased complexity also requires connectivity to the outside world and more processing power and storage demands for the collected data.

The logical escape was to use the power of cloud data centers, so the next demand was to connect IoT sensors and controller devices to the Internet and cloud servers, as Figure 1 shows.

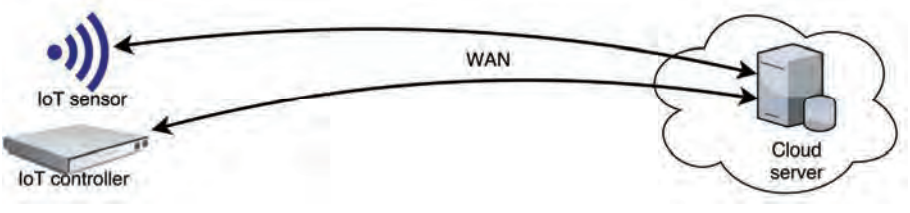

Figure 1. Requirements to connect IoT sensor and/or controller device to the cloud.

Connection to the cloud is solved rather efficiently in two ways: by Internet providers and mobile operators. The first scenario is realized by use of the Internet and local networks. End-users of IoT devices can use their Internet connection and the LAN installed at their premises. To connect to the IoT sensors, they set up a WiFi router or use a direct LAN cable connection. Figure 2 presents this idea of establishing a connection to the cloud server.

Figure 2. Requirements to set up a WiFi to connect the IoT sensor and controller device to the cloud server

This solution works efficiently while the amount of transmitted data is not too high and does not cause communication bottleneck. However, the design of IoT sensors and controller devices may use streaming data with big volume and velocity.

Another feature in the development of IoT sensors and devices was to make them mobile and wireless, free of any obstacles while moving. This means that they cannot be connected to a static electrical power supply, but, instead, that they require small batteries. So, a problem arises due to limited resources because the processing should be small enough to conserve energy and, at the same time, enable longer battery life prior to recharging.

Limited resources do not only constrain the processing, but also the use of wireless connections. The solution is to use small personal area network connections to conserve the resources of the IoT sensor and devices for as long as possible. Typical examples include Bluetooth or other related wireless network technologies. This changed the overall concept because WiFi requires a more powerful wireless connection and is not appropriate for battery-operated devices.

The solution to this problem was to create new devices with a special purpose to accompany IoT devices and function as digital signal repeaters. These devices are connected to the IoT sensor or controller device via a personal area network (such as a direct cable or Bluetooth) and can communicate with the WiFi router, as Figure 3 shows. Their only function is to receive small power signals, buffer them and then transmit to the destination over WiFi or other LAN connection.

Figure 3. IoT sensor and controller devices connected to a cloud server via a digital repeater and WiFi connections.

Some of these devices are no longer required to be mobile or wirelessly connected, but they can have a regular electrical power supply and be connected to the LAN installed at the premises where they are hosted. The other will need to recharge but their size can allow a higher battery capacity and decrease the need for frequent recharging. A typical example of a device in this category is a smartphone, which can communicate to the personal wearable sensors from one side and the WiFi from the other side.

The solution presented in Figure 3 is a typical solution of an IoT sensor or controller device to the cloud using a standard Internet connection.

Cloud-based solutions

We next explain the alternative to a direct Internet connection. Mobile operators offered a solution to this problem—instead of setting a WiFi router and Internet connection, they offered a mobile network solution. To come closer to IoT sensors, they use 3G/4G or 5G connections from their operated services. This solution requires a new device capable of accepting signals from the personal area network and connecting to the mobile operator's network. The user can use a classical mobile device and connect to the IoT devices by Bluetooth. The market also offered solutions with special dedicated devices that can communicate with Bluetooth and also accept a SIM card to connect to the mobile operator's network. Figure 4 presents a typical solution provided by the mobile operator for the purpose of IoT sensors and remote controller devices.

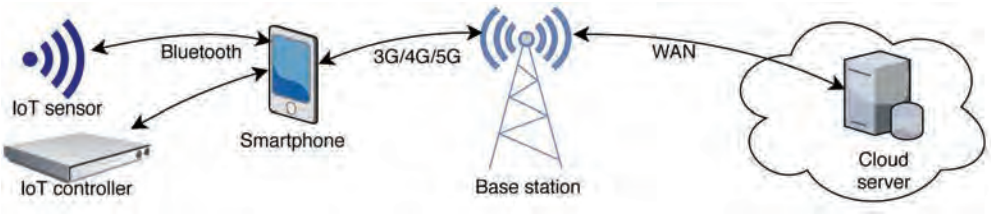

Figure 4. IoT sensor and controller devices connected to a cloud server via smartphone and mobile operator's network.

These designs are good enough if the IoT sensors do not require a lot of communication and processing. For example, in the case of measuring environmental temperature, there is no need to transmit data permanently, but at regular intervals, such as one value per minute, or less.

New technologies enabled more sophisticated sensors that permanently sense data with a high frequency. For example, a camera can sense the traffic density or an ECG sensor can sense heartbeats. The birth of big data is based on such streaming sources that provide data with high volume, velocity, variety, variability, and veracity.

The big data evolution introduced new functions in the IoT world and characteristics for IoT devices, and also addressed the need for increased storage and processing requirements. The previous design considerations, now introduce new problems such as where to store the streaming data and who will process it to trigger relevant activities on the controller device.

The existing designs of cloud-based IoT (Figure 3) or mobile operated cloud-based IoT solution (Figure 4) are not good enough. The problems occur because streaming data transmitted over the WAN will soon occupy the network throughput as the number of these connections increases and provokes server congestion. Of course, there is a solution for the increased WAN bandwidth, but the number of connected IoT sensors rises with a higher rate than the network providers can cope with.

Cloudlet and mobile edge computing

The solution for avoiding communication bottlenecks is to distribute the processing closer to the user. This is the basis of a cloudlet solution. A smaller server is hosted closer to the user, and all the storage and processing requirements are now distributed to these smaller cloudlet servers. They provide the same functions as the cloud data center, but can accept only a limited number of IoT devices and provide relevant service. The cloud data center is now relieved with high communication demands and collaborates with the cloudlet to exchange relevant information and provide extended services with added value.

Figure 5 presents the architecture of the cloudlet solution for streaming IoT devices. The IoT sensor or controller devices communicate with the digital repeater via a personal area network, using wireless Bluetooth technology or a direct cable connection. The digital repeater buffers the received signal and then transmits it to the cloudlet server via a WiFi router. The cloudlet server

collects the streaming data and processes it, providing relevant services. Only selected and processed information is sent to the cloud data center, which provides a more comprehensive analysis of the input data. This solution is a successful upgrade of the direct cloud solution due to the smaller latencies and the distribution of high throughput streaming data and processing to localized servers. It has been proven that the latencies generated by the digital repeater and WiFi router, the Bluetooth/WiFi and LAN connections, and slower cloudlet servers are smaller than the overall latencies to access the cloud data center even if it processes faster than the cloudlet server. The highest impact of this design is the decrease of high throughput demands.

Figure 5. Cloudlet solution for streaming IoT sensor and controller device.

Mobile operators have offered a similar concept to solve the high throughput demands of IoT devices (Figure 6). They host smaller servers at the premises of their base stations, and bring the storage and processing closer to the users at the edge of their network. This is known as the mobile edge computing solution.

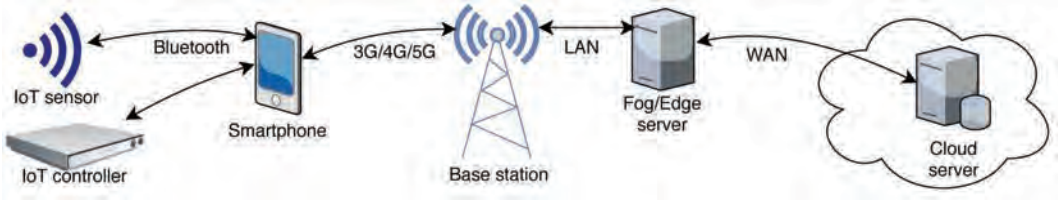

Figure 6. Mobile edge computing solution for streaming IoT sensor and controller device.

Edge computing solutions

So far, we have presented four solutions for the IoT world's requirements for computation and storage offload:

- Cloud directly connected via Internet (Figure 3).
- Mobile operated cloud solution (Figure 4).
- Cloudlet solution (Figure 5).
- Mobile edge computing solution (Figure 6).

The first and third use a kind of open environment since the user can select the Internet provider and cloud provider. The second and last solutions are proprietary for mobile operators and the user is usually in a "lock-in" situation with no possibility of choosing or changing anything. In the cloudlet solution, the user needs to install his or her own cloudlet server and pay a monthly subscription fee. In the case of mobile operator solutions, the monthly-prepaid subscription is the only alternative.

The last two designs, which solve the problem of streaming IoT data, encompass the movement of centralized services closer to the user (an implementation of the edge computing concept to push applications, data, and computing power by bringing the processing closer to the logical extremes). The architecture of cloudlets or the mobile edge computing solution assumes the establishment of an intermediate layer between the cloud server and the IoT device in the form of an edge server.

This leads back to the concept of moving the data collection and initial processing even closer to the user and closer to the place where the data is produced and used. This concept means that there will be another layer, an additional edge device, between the end-user IoT device and the intermediate layer.

The idea presented in Figure 7 changes the use of the repeater with a smaller server in a form of an edge device (sometimes called dew[9]). It can be a desktop computer, laptop, tablet, smartphone or another mobile device capable of processing the incoming data stream locally. In addition, it can store data for a limited time, and unless the data is unloaded, the old data will be deleted. It can communicate both to the lower hierarchy level (IoT devices) and higher levels (cloudlet or mobile edge computing server and cloud data center).

Figure 7. Edge computing solution using an IoT and edge devices.

The communication with IoT devices consists of reception of data and transmission of triggers. The higher levels enable interconnection and information exchange with the outside world, for example, to unload the stored data and send relevant information to the cloud data center. This edge computing design brings the data collection and processing closer to the source and end-users.

DISCUSSION

Dolui and Datta give a good comparison of related terms when they compare terms based on their implementations.[10] They analyze fog computing node devices by routers, switches, or gateways, with cloudlets as data centers in the nearby LAN, and mobile edge computing as servers at the base stations of mobile operators.

Edge computing has been analyzed from various architectural approaches. Nastic et al. introduced serverless architecture for edge computing, so the idea of going closer to the roots of data producers is even more extracted.[11]

Analysis of the best edge computing solution

Our idea was to choose the best edge computing solution to distribute processing closer to the user and data source. One would ask which solution is the best, cloudlet, fog or mobile edge computing, or a combination of these concepts. There is no straight answer to this question, so we will analyze the conditions that help solution providers choose the best.

The direct cloud solution opposed to the edge computing concept is the best alternative in the case of a static IoT sensor that senses environmental data that does not change a lot over time, such as temperature. This sensor does not stream data but rather sends information on regular time intervals and, in this case, the communication, storage and processing requirements are very small. This is especially important in the case of a regular power supply of electricity.

The edge computing solution with a cloudlet server and cloud data center is the best in the case of a static IoT sensor that streams a lot of data, such as video signal. The sensor transmits a lot of data that is stored for a limited time period on the edge device and selected data is sent to the cloudlet for further processing in collaboration with the cloud data center.

The mobile edge computing solution is a good alternative if the sensor is not static (it changes its geographic position) and there is no available Internet provider coverage. In this case, the sensor

directly communicates with the mobile operator's network where the fog/edge server will collect and process incoming data, collaborating with their cloud data center.

The mobile edge computing and cloudlet solutions are the best alternative for a movable and streaming IoT sensor and edge device. The mobile edge device may not have a continuous Internet connection. Since this edge device collects, processes and stores data, it is not necessary to have a permanent Internet connection. In this case, whenever the IoT sensor and edge device are within the range of an Internet connection, the data will be unloaded and the cloudlet/cloud server will update the current situation. It is also possible to use a 3G/4G connection to exchange data with the cloudlet/cloud server.

Table 1 gives a summary of the best solutions in various scenarios for the IoT sensor and its corresponding edge device. Based on their position, the following categories are determined:

- static location, when the IoT sensor and edge device do not move over time and have permanent energy supply;
- movable location in the case of a movable IoT sensor; and
- edge device, charged by a battery.

In our classification, the sensor streams data if the sampling frequency is higher than 1Hz at one data item per second. We understand that this value is not determined exactly, but usual values for streaming data are those with sampling frequencies higher than 100 Hz and those that transmit data less than 0.1 Hz are assumed not to be streaming.

Table 1. The Best solution in the case of static and movable IoT sensors based on data streaming and Internet availability.

Position/ Data	Streaming Data	Sending Files
Static	Edge device and cloudlet/cloud	Cloud only
Movable	Edge device and MEC	Fog only
Movable with WiFi	Edge device and cloudlet/cloud	Cloud only

In the presented categorization, we also differ between a permanent Internet connection and availability of mobile-operated network access.

A CASE STUDY: A WEARABLE ECG SENSOR

Let's analyze the ICT requirements of an existing health-related IoT solution that uses an edge device. The case study consists of a wearable ECG sensor attached to the human body and a mobile phone that acts as an edge device. The communication between the wearable sensor and the edge device is via Bluetooth and via WiFi or 3G/4G for the Internet.

The users prefer to wear their ECG sensors and monitor heart functions on the edge device (smart-phone) without any constraints that prevent their normal physical activities. They prefer mobile phones whose batteries last more than 20 hours without recharging since it is common practice to recharge the mobile phone at night while the user is sleeping.

Storage requirements

The edge device (smartphone) stores data with a limited capacity. The assumption is that the attached wearable ECG sensor samples data with a 500 Hz frequency and uses a 2-byte sample. So the edge device (smartphone) needs to store 1KB per second, or 86.4 MB per day. Unless the edge device (smartphone) does not unload and delete old data, the requirement to store the incoming data will reach 2.6 GB in one month.

Annual storage requirements will exceed 949 GB closing to the TB limits.

Communication requirements

If all data is streamed over the WiFi to higher level servers, then the demands of this system require 8Kbps transfer rate if streaming data is transmitted without identification, time, and space stamps. However, application software may involve a big overhead, for example, complex .Net methods transmit textual data instead of binary numbers.

A typical TCP/IP connection usually adds overhead due to the headers of the implemented TCP/IP levels. Considering the acknowledgements and other controlling mechanisms, we can expect this transfer rate to double.

Even a doubled transfer rate value is acceptable since the current standard cable operators offer uplink speed of at least 1Mbps. So, the communication with the cloud will not cause a substantial problem from the user's perspective.

However, a communication bottleneck may appear at the cloud data center if hundreds or thousands of users log into the cloud and concurrently stream data. The problem arises mostly at the server side, at the premises of the cloud data center.

Processing requirements

The processing requirements depend mostly on the implemented algorithms. Small pipelined ECG analysis and feature extraction algorithms may require up to 1K computations per sample, but comprehensive analysis may require more than 1M operations per sample.

Assuming that 500 samples arrive each second then the requirement will be 500K operations per second for streaming pipelined algorithms and 500M operations per second for comprehensive analysis. Theoretically, all modern smartphones can perform sufficient operations per second, for example, iPhone 6 reaches 2 Gflops/s, and Samsung Galaxy S6 even higher.

However, if hundreds or thousands of ECG sensors stream data to a cloud center, then a bottleneck should be expected at the premises of the cloud data center.

Energy consumption

So far, we have analyzed the ICT requirements and concluded that the edge device (smartphone) can provide sufficient capacity to cope with them. However, the real problems start with an analysis of energy consumption.

Smartphone producers claim that the battery of their models from 2016 have a capacity of at least 1.800 mAh and can play video for at least 10 hours until they run out of battery. However, in terms of 3D gaming, they can only survive about two-and-half hours. Both these examples show that in case of increased processor activity, the battery life is not sufficient for a careless 24 hour use. A careful analysis shows that most of the energy is consumed by the screen, which is not required in the case of a monitoring application. The screen will be used only if the user likes to monitor the activities, or in the case of an emergency. Several tests show that the transmission of the ECG signal via WiFi or 3G/4G can be realized with available battery capacity. However, the problem is the software solution and the implemented algorithm to detect the heartbeat and diagnose the heart function. An algorithm using a comprehensive analysis of the ECG signal will require too much processor activity that will drain the battery life in at least 4 hours (less than 3D video gaming). So there must be a compromise, and the developers have to code software with highly efficient streaming algorithms.

Recommendations

Wearable IoT (ECG) sensors and edge devices (smart-phones) can be built in such a way that enables sufficient capacity for the communication, storage and processing requirements. However, the real problem arises when analyzing energy consumption of the wearable edge device.

The compromise is to decrease some of the analyzed communication, storage, and processing requirements. The first solution is to select what is to be transmitted to the cloud, the second is to select what data to store (not all data), and the last is to realize less precise monitoring algorithms.

Usually, the solution providers compromise with a combination of all these requirements or select the last one: to develop smaller monitoring capacities and transfer all data to the cloud, where a comprehensive analysis can be realized. This solution is designed to remain on the safe side, keep and transmit all data to the server, and enable essential monitoring information.

In the case of the ECG, some other relevant issues can also arrise, since not all data is needed unless the heart is detected to be out of normal function. This is why the edge device will play a huge role in further development of this computing concept. Only selected data will be transmitted and stored on the cloud. This is mostly comprised of the problematic issues since data that represents normal cardiac activity is not intended for further analysis. Of course, a limited storage capacity of the cloud can be offered to store a certain time interval, such as one or past two months. The old data will be deleted unless it is marked to be saved.

CONCLUSION

In this paper, we have presented the evolution of the edge computing concept from an IoT perspective. In the past, IoT sensors and controller devices required processing that is done by connection to a cloud data center, either via WiFi or via 3G/4G or 5G. This concept successfully completes centralized processing in the cloud environment.

However, the evolution of IoT sensors with streaming data, changed the course of events. At the moment of writing communication and processing are brought closer to the user, either by establishing a cloudlet server or by a mobile edge computing server on the edge of the mobile operator network. The new edge computing concepts go a step beyond this concept, bringing the communication and processing even closer to the user, and hosting an edge device close to the IoT devices and end users. This evolution showcases how computing architecture has changed over time, first we had a centralized solution, and now we are going back to distributed end-user devices.

Finally, edge computing enables independence of constant networking, reducing processing latency and ,at the same time, allows autonomous functioning and collaboration with distant servers.

REFERENCES

1. J. Gubbi et al., “Internet of things (IoT): A vision, architectural elements, and future directions,” *Future Generation Computer Systems*, vol. 29, no. 7, 2013, pp. 1645–1669.
2. W. Shi and S. Dustdar, “The Promise of Edge Computing,” *Computer*, vol. 49, no. 5, 2016, pp. 78–81.
3. I. Stojmenovic and S. Wen, “The fog computing paradigm: Scenarios and security issues,” *2014 Federated Conference on Computer Science and Information Systems* (FedCSIS 14), 2014, pp. 1–8.
4. F. Bonomi et al., “Fog computing and its role in the internet of things,” *Proceedings of the first edition of the MCC workshop on Mobile cloud computing*, 2012, pp. 13–16.
5. M. Satyanarayanan et al., “The case for VM-based cloudlets in mobile computing,” *IEEE Pervasive Computing*, vol. 8, no. 4, 2009, pp. 14–23.

6. A. Bahtovski and M. Gusev, "Cloudlet challenges," *Procedia Engineering*, vol. 69, 2014, pp. 704–711.
7. M.T. Beck et al., "Mobile edge computing: A taxonomy," *The 6th International Conference on Advances in Future Internet* (IARIA 14), vol. IARIA 14, 2014, pp. 48–54.
8. P. Mach and Z. Becvar, "Mobile edge computing: A survey on architecture and computation offloading," *IEEE Communications Surveys & Tutorials*, 2017.
9. M. Gusev, "A dew computing solution for IoT streaming devices," *40th IEEE International Convention on Information and Communication Technology, Electronics and Microelectronics*, vol. MIPRO 17, 2017, pp. 387–392.
10. K. Dolui and S.K. Datta, "Comparison of edge computing implementations: Fog computing, cloudlet and mobile edge computing," *Global Internet of Things Summit* (GIoTS 17), 2017, pp. 1–6.
11. S. Nastic et al., "A serverless real-time data analytics platform for edge computing," *IEEE Internet Computing*, vol. 21, no. 4, 2017, pp. 64–71.

ABOUT THE AUTHORS

Marjan Gushev is a computer scientist and professor at University Sts. Cyril and Methodius Faculty of Information Sciences and Computer Engineering. His research interests include parallel processing, computer networks, and Internet technologies. Gushev received a doctoral degree in computer science from the University of Ljubljana, Slovenia. Contact him at marjan.gushev@finki.ukim.mk.

Schahram Dustdar is a full professor of computer science, and he heads the DSG at TU Wien. His work focuses on distributed systems. Dustdar is an IEEE Fellow, a member of the Academia Europaea, an ACM Distinguished Scientist, and recipient of the IBM Faculty Award. He is on the editorial boards of IEEE Internet Computing and Computer. He's an associate editor of IEEE Transactions on Services Computing, ACM Transactions on the Web, and ACM Transactions on Internet Technology. He's the editor-in-chief of Springer Computing. Contact him at dustdar@dsg.tuwien.ac.at; http://dsg.tuwien.ac.at/staff/sd/.

This article originally appeared in
IEEE Internet Computing, *vol. 22, no. 2, 2018.*

EDITOR **MARK CAMPBELL**
Trace3; MCampbell@trace3.com

Smart Edge: The Effects of Shifting the Center of Data Gravity Out of the Cloud

Mark Campbell, Trace3

A confluence of macrotrends is pushing data, computing, and analytics into the rapidly expanding world at the edge of the connected ecosystem. What is shifting, what is causing it, and what are the overall effects?

During the past half century, the IT industry has witnessed several pendulum swings that moved the center of gravity from centralized workloads to decentralized processing and back. Disparate analog and mechanical devices consolidated into centralized digital mainframes then spread into isolated personal computers, came together to form interconnected client/server processing and cloud services, and have now separated into the far-flung devices that comprise the Internet of Things (IoT).

As early as 1994, the advent of embedded microprocessors necessitated the definition of a topology to describe and manage the connected ecosystem.[1]

- *Cloud*: Also referred to as the core, the cloud is the collection of fixed, interconnected data centers that house infrastructure, applications, and data.
- *Edge*: The edge is the collection of remote devices that couple sensors, processors, storage, and a network connection to the cloud. Referred to collectively as the IoT, these devices are often mobile and, thus, use wireless communications, have low power consumption, and provide only minimal processing and storage capacity.
- *Fog*: This is the hierarchical interconnection between the edge and cloud, including localized processing, gateways, networking components, access devices, and security controls.

Digital Object Identifier 10.1109/MC.2019.2948248
Date of current version: 22 November 2019

To say that the IoT is exploding is an understatement. It is estimated there are five IoT devices for every person on the planet, with an annual growth rate of 12%, and there will be 125 billion of them by 2030.[2] Outpacing the growth in the number of devices is the meteoric rise in IoT revenue, which is expected to exceed US$1.7 trillion in 2019, up 350% from US$486 billion in 2013.[3] This expanding galaxy of connected devices produces an astronomic amount of information and shifts the center of data gravity from the cloud to the edge. By 2025, it is estimated that more data will be created in IoT devices than in data centers: 90 ZB out of a global datasphere of 175 ZB.[4]

> The reigning data processing philosophy is to save all of yesterday's data today, since it is unknown what could be important tomorrow.

CHALLENGES

As data sets grow, they become proportionally difficult to move. This causes more transportable consumers of the data, such as applications and services, to gravitate toward the source to decrease latency and increase throughput.[5] As on-device data burgeon, the center of data gravity shifts out of the cloud and creates many problems at the edge, including the following:

- *Bandwidth*: The amount of data created at the edge will soon outstrip the connectivity bandwidth to ship the information to the cloud. For example, it is estimated that autonomous cars will generate between 5 and 20 TB per vehicle per day.[6] At those rates, even the 10-Mb/s bandwidth offered by 5G wireless communications technology will quickly be consumed.
- *Latency*: Many of today's distributed IoT applications require cloud processing. As the center of data gravity moves toward the edge, the access latency between cloud-hosted applications and device-based data will increase, creating serious problems for systems that require real-time responses, such as manufacturing, health care, security, and power distribution.
- *Outages*: The edge depends on connectivity. An outage in a service, server, or network could have serious implications for edge systems that rely on other resources for processing, data, security, and command and control. It is also difficult to predict how a local outage could propagate across other devices and whole segments of the edge in a cascading failure.
- *Security*: As data move to the edge, the attack surface of a given enterprise increases exponentially. Each edge device, IoT link, distributed application, and edge user becomes another potential target for exploitation by bad actors.
- *Privacy*: Much of the data and metadata generated by IoT devices contains personal and privileged information. As data move to the edge, questions of ownership, rights, and privacy become increasingly complex.

To help lessen the impact of these problems, computational resources are migrating toward the edge and interconnecting fog infrastructure. Commonly called *fog computing* or *cloudlets*, small-scale mini data centers located near the edge provide resource-intensive and transaction-heavy edge applications with powerful computing resources and lower latency. As this shifting data gravity pulls computing to the edge and fog layers, the impact of bandwidth, latency, and outage issues is lessened but not eliminated.

With the increasing processing horsepower outside the cloud, advanced analytic applications are also moving toward the edge, bringing with them their embedded artificial intelligence (AI) models. By performing analytics at or near the edge, decision making is pushed closer to both the sensors providing telemetry data input and the actuators reacting to analytics-based output. Additionally, emerging parallel programming paradigms, such as OpenACC and the CUDA-aware message passing interface, are changing edge applications to decrease latency and gain a degree of autonomy in the event of connection outages between the edge and the cloud.

THE SMART EDGE

These transitions have triggered the emergence of the "smart edge," which is composed of edge data, computing, analytics, and AI. Smart edge solutions are being applied across numerous industries:

- *Health care*: Health care is embracing the rapid adoption of smart edge medical devices, including wearables, hearables, injectables, and ingestibles for patients and computer vision, remote patient monitoring, and telehealth for clinicians.[7] The evolution of medical edge devices mirrors the fog computing developments in the smart hospital, where connected medical device data are fed to advanced analytical platforms to derive insights that are accessible to local and remote users and systems.
- *Agriculture*: A bevy of smart sensors and processors is fueling a surge in smart farming, where

smart edge devices are used for video-based automation, autonomous seeding/feeding/weeding/reaping, and real-time telemetry from drones and robotic machinery. The smart edge provides more accurate and insightful environmental monitoring of soil, crop, weather, and chemical data. It is also a prerequisite for agro-innovations, such as vertical farming, which uses sensor data for local analytics to optimize food production without the need for farmland.[8]

- *Transportation*: The self-driving car has become the de facto mascot for smart edge technology, and the obligatory references to lidar processing, data uplink issues, and security risks are almost cliché. However, the smart edge is not only the catalyst for advancements in consumer transportation but it enables autonomous shipping, rail, trucking, and air services and, by extension, smart ports, roads, and traffic. As exciting as these innovations are, the real power that the smart edge unlocks for transportation is in vehicle-to-everything (V2X) technology, where vehicles and traffic systems all communicate directly (that is, bypassing the cloud) through industry standards such as IEEE 802.11p/Dedicated Short-Range Communications, the 3rd Generation Partnership Project-defined cellular V2X, and 5G. V2X promises safe, agile, and efficient autonomous transportation across all vehicle categories.[9]
- *Fog computing and micro data centers*: As smart edge developments enable use cases in all industries, there is an implicit dependency on a smart fog computing hierarchy that collects and distributes data, users, policies, and security across the edge ecosystem. As with the smart edge, fog computing leverages AI to place computing, storage, network, and security resources closer to the points of creation and consumption. A compelling trend is the packaging of fog computing resources in micro data centers, which consolidate resilient, high-capacity, low-touch infrastructure in a refrigerator-size cabinet located locally and linked hierarchically. Micro data centers are quietly appearing in various distributed outlets, such as bank branches, manufacturing facilities, retail outlets, smart buildings, telecommunications, and the power grid.[10]

The advent and maturation of the smart edge creates many subtle consequences that are only now being explored. What happens when smart edge analytic demands exceed the capabilities of the edge device? What happens when edge data are too large to uplink to the fog and cloud layers? How will direct edge-to-edge communication solve the problems that it shares with the edge-to-cloud system (outages, bandwidth, latency, and security)? How will smart edge devices secure their rapid-fire chatter against eavesdroppers when encryption bloats compressed data and amplifies the effects of data gravity? How will the smart edge provide "explainability" to trace and remediate local and cascading faults when smart apps typically present only a black-box cognitive model?

THE LEADING EDGE

While these and other challenges emerge, the overwhelming promise of the smart edge is compelling venture capital firms to fund and deploy start-ups to exploit the "edge effects." General IoT investment rose to US$16.7 billion in 2018, a 94% growth rate from 2017.[11] The adventurous start-ups are racing to provide solutions to a variety of smart edge opportunities, including the following:

- *Stratified learning*: Today's monolithic learning pipelines train AI models in the cloud and download them to the fog and edge layers. As feedback data trickle back to the cloud from the edge, the model is retrained and repropagated out. With smart edge and fog computing, this cycle is trifurcated with rapid localized and regionalized learning that occurs as close to the data creation and sensor/actuator interaction as possible. The resulting locally trained models filter back to the cloud, where they are amalgamated into the master model.[12] Although it is still experimental, the implications of real-time, local learning for high-volume streaming edge data, especially high-resolution video, will reshape many industries, such as manufacturing, defense, security, transportation, and surveillance.
- *Smart data digest*: Whether a smart edge device is a phone, hydroelectric generator vibration monitor, surgical telehealth link, or industrial manufacturing robot, there is a wealth of data volume but a sparsity of data information. The reigning data processing philosophy is to save all of yesterday's data today, since it is unknown what could be important tomorrow. While logically sound, this approach is empirically sloppy. Emerging smart data digest and data reduction solutions attempt to find meaning-rich needles and discard the meaningless haystacks. Although this seems akin to picking next week's lottery numbers, emerging techniques and solutions that use observation and attribute reduction, digital threading, and principal component analysis

This article originally appeared in Computer, *vol. 52, no. 12, 2019.*

are being developed to target edge data reduction.[13]

- *Personal analytical sovereignty*: Data sovereignty provisions, including the European Union's General Data Protection Regulation, govern which country's laws have jurisdiction over a given data set. Fueled by events such as the Cambridge Analytica debacle, personal data sovereignty measures, such as the California Consumer Privacy Act, decide who should control the collection, dissemination, and deletion of private information. Personal analytical sovereignty takes this a step further by determining who controls the AI models that are trained from personal data. Why the distinction? If the sovereignty of personal data on a smart edge device is determined to reside with the user, other interested parties can simply bypass personal data protection and use the much more valuable AI models. For instance, does my personal retail buying model belong to me, my phone company, the store I am visiting, or the edge e-commerce software I installed that eagerly learns my buying habits? Solutions that resolve these issues cleanly, which is by no means an easy feat, can grant their users a portable, secure, and personally controlled behavioral model that may share the personal data needed for transactions but will retain and control the learned persona of its owner.

The IoT data explosion has created an ever-increasing volume of edge data, and the resulting edge data gravity is causing significant challenges in today's connected ecosystem. In response, compute capacity and AI-based analytics resources are drawn closer to the edge, creating the smart edge. While smart edge devices and fog computing have made considerable inroads into all industries, manifold issues have emerged, including bandwidth, latency, security, and privacy. However, each of these problems leads to the next generation of smart edge solutions and opportunities in a regenerative feedback cycle. Stratified learning will enable real-time cognitive processing and autonomous training, while the smart data digest will preserve bandwidth by increasing the data value and obviating many of the effects of edge data gravity. However, the greatest smart edge impact may be the realization of hyperadaptive individualized computing that is protected by personal analytical sovereignty. We are just seeing the beginning.

REFERENCES

1. R. S. Raji, "Smart networks for control," *IEEE Spectr.*, vol. 31, no. 6, pp. 49–55, 1994.
2. J. Howell, "Number of connected IoT devices will surge to 125 billion by 2030, IHS Markit says," IHS Markit Technology, Oct. 24, 2017. [Online]. Available: https://technology.ihs.com/596542/number-of-connected-iot-devices-will-surge-to-125-billion-by-2030-ihs-markit-says
3. CB Insights, "What is edge computing?" CB Insights Research, New York, Aug. 8, 2018. [Online]. Available: https://www.cbinsights.com/research/what-is-edge-computing/
4. D. Reinsel, J. Gantz, and J. Rydning, "The digitization of the world from edge to core," IDC, Framingham, MA, White Paper, 2018.
5. D. McCrory, "Data gravity: In the clouds," Data Gravitas, United Kingdom, 2010.
6. M. Pastor, "Driverless cars generate massive amounts of data. Are we ready?" Enterprise AI, San Ramon, CA, Oct. 24, 2018. [Online]. Available: https://www.enterpriseai.news/2018/10/24/driverless-cars-generate-massive-amounts-of-data-are-we-ready/.
7. Sciforce, "IoT in healthcare: Are we witnessing a new revolution?" Medium, 2019. [Online]. Available: https://medium.com/sciforce/iot-in-healthcare-are-we-witnessing-a-new-revolution-6bb0ecf55991
8. Lanner, "5 Edge computing use cases for smart farming and agriculture," Lanner America, Oct. 1, 2018. [Online]. Available: https://www.lanner-america.com/blog/5-edge-computing-use-cases-smart-farming-agriculture/
9. Research and Markets, "Global V2X (vehicle-to-everything) communications ecosystem markets, 2019–2030: By 2022, V2X will account for a market worth $1.2 billion," Res. Markets, Dublin, Ireland, Rep. 4759721, 2019.
10. L. Dignan, "What's next for data centers? Think micro data centers," ZDNet, Apr. 14, 2019. [Online]. Available: https://www.zdnet.com/article/whats-next-for-data-centers-think-micro-data-centers/
11. L. Columbus, "Top 25 IoT startups to watch in 2019," *Forbes*, Feb. 3, 2019. [Online]. Available: https://www.forbes.com/sites/louiscolumbus/2019/02/03/top-25-iot-startups-to-watch-in-2019/#4baf81503cc0
12. Z. Doffman, "Network effects: In 2019 IoT and 5G will push AI to the very edge," *Forbes*, Dec. 28, 2018. [Online]. Available: https://www.forbes.com/sites/zakdoffman/2018/12/28/network-effects-in-2019-iot-and-5g-will-push-ai-to-the-very-edge/#3ab09e3f6bbe
13. Karwan, "Robust & non-robust methods for data reduction that you should know," Medium, Aug. 27, 2019. [Online]. Available: https://towardsdatascience.com/robust-non-robust-methods-for-data-reduction-you-should-know-a1dc7347a802

MARK CAMPBELL is the chief innovation officer for Trace3. Contact him at MCampbell@trace3.com.

Keep up with the latest IEEE Computer Society publications and activities wherever you are.

Follow us:

@ComputerSociety

facebook.com/IEEEComputerSociety

IEEE Computer Society

youtube.com/ieeecomputersociety

instagram.com/ieee_computer_society

IEEE COMPUTER SOCIETY

CALL FOR PAPERS

IEEE 21st International Conference on Information Reuse and Integration for Data Science (IEEE IRI 2020)

Las Vegas, NV, USA / 11 – 13 August 2020

The IEEE IRI conference is a recognized forum for researchers and practitioners from academia, industry, and government to present and exchange ideas that address real-world problems with real-world solutions.

This conference is now seeking excellent, novel, and contemporary papers covering all aspects of Data—including Scientific Theory and Technology-Based Applications.

The conference includes, but is not limited to, the areas listed below:

- Application—Autonomous Vehicles, Business, Education, Engineering, Healthcare, the Internet of Things, Math, Military, Multimedia, NLP, Robotics, Science, Security, Social Networking, Space, Vision, et al.
- Contemporary as well as Novel Data Mining Techniques
- Data & Knowledge Representation and Management
- Data Science & Technologies—Heuristic Acquisition
- Data Visualization
- Graph Models
- Machine Learning & AI
- Predictive Data Analysis & Intelligence
- Predictive Modeling
- Recommender Systems
- Statistical Analysis
- Theory

This year, for the first time, IRI will hold classified military sessions at the nearby Naval Information Warfare Center (NIWC) in San Diego. Presentation of classified IRI military papers and attendance at these sessions is open to those holding an active US government Secret clearance. A clearance is not required for any other IRI sessions. SECRET presentations pertaining to the following areas of interest are encouraged:

- Artificial Intelligence, Heuristics, and Explanation-Based Learning
- Machine Learning
- Computer Vision
- Hypersonic Flight
- Autonomous Vehicles
- Predictive Maintenance
- Logistics
- Silicon Compilers
- Quantum Theory and Application

www.bit.ly/iri20-cfp

Important Deadlines

Abstract submission: 15 April 2020

Full paper research/industry/application/gov't track deadline: 22 April 2020

Short paper track deadline: 22 April 2020

Poster and demo paper track deadline: 8 May 2020

Full/short paper acceptance notification: 1 June 2020

Poster/demo paper acceptance notification: 15 June 2020

Camera ready submission deadline: 20 June 2020

Author registration due: 1 July 2020

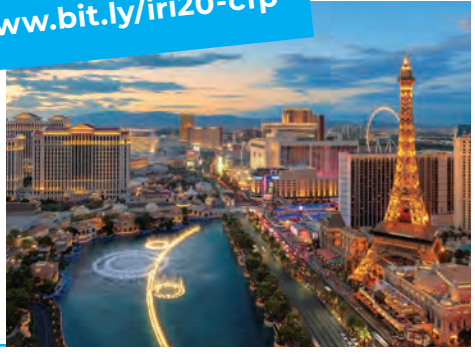

IEEE COMPUTER SOCIETY

IEEE

Get Published in the New *IEEE Open Journal of the Computer Society*

Submit a paper today to the premier new open access journal in computing and information technology.

Your research will benefit from the IEEE marketing launch and 5 million unique monthly users of the IEEE *Xplore*® Digital Library. Plus, this journal is fully open and compliant with funder mandates, including Plan S.

Submit your paper today!
Visit www.computer.org/oj to learn more.

IEEE COMPUTER SOCIETY | IEEE

IEEE INTERNATIONAL SYMPOSIUM ON HARDWARE-ORIENTED SECURITY AND TRUST

4–7 May 2020 • San Jose, CA, USA • DoubleTree by Hilton

Join dedicated professionals at the IEEE International Symposium on Hardware Oriented Security and Trust (HOST) for an in-depth look into hardware-based security research and development.

Key Topics:

- Semiconductor design, test and failure analysis
- Computer architecture
- Systems security
- Cryptography and cryptanalysis
- Imaging and microscopy

Discover innovations from outside your sphere of influence at HOST. Learn about new research that is critical to your future projects. Meet face-to-face with researchers and experts for inspiration, solutions, and practical ideas you can put to use immediately.

REGISTER NOW: www.hostsymposium.org

IEEE COMPUTER SOCIETY · IEEE · tttc · IEEE TC on Security & Privacy

Conference Calendar

Questions? Contact conferences@computer.org

IEEE Computer Society conferences are valuable forums for learning on broad and dynamically shifting topics from within the computing profession. With over 200 conferences featuring leading experts and thought leaders, we have an event that is right for you.

Find a region: Africa ■ Asia ▲ Australia ◆ Europe ● North America ◗ South America ★

March

2 March
- WACV (IEEE Winter Conf. on Applications of Computer Vision) ◗

9 March
- DATE (Design, Automation & Test in Europe Conf. & Exhibition) ●
- IRC (IEEE Int'l Conf. on Robotic Computing) ▲

16 March
- ICSA (IEEE Int'l Conf. on Software Architecture) ★

22 March
- VR (IEEE Conf. on Virtual Reality and 3D User Interfaces) ◗

23 March
- ICST (IEEE Conf. on Software Testing, Validation and Verification) ●
- PerCom (IEEE Int'l Conf. on Pervasive Computing and Communications) ◗

April

5 April
- ISPASS (Int'l Symposium on Performance Analysis of Systems and Software) ◗

9 April
- MIPR (IEEE Int'l Conf. on Multimedia Information Processing and Retrieval) ▲

13 April
- Mobile Cloud (IEEE Int'l Conf. on Mobile Cloud Computing, Services, and Eng.) ●
- SOSE (IEEE Int'l Conf. on Service-Oriented System Eng.) ●
- DAPPS (IEEE Int'l Conf. on Decentralized Applications and Infrastructures) ●
- JCC (IEEE Int'l Conf. on Joint Cloud Computing) ●
- AITest (IEEE Int'l Conf. on Artificial Intelligence Testing) ●
- BigDataService (IEEE Int'l Conf. on Big Data Computing Service and Machine Learning Applications) ●

14 April
- PacificVis (IEEE Pacific Visualization Symposium) ▲

15 April
- COOL Chips (IEEE Symposium on Low-Power and High-Speed Chips and Systems) ▲

20 April
- ICDE (IEEE Int'l Conf. on Data Eng.) ◗

May

3 May
- FCCM (IEEE Int'l Symposium on Field-Programmable Custom Computing Machines) ◗

Published by the IEEE Computer Society

4 May

- HOST (IEEE Int'l Symposium on Hardware Oriented Security and Trust) ◗

18 May

- SP (IEEE Symposium on Security and Privacy) ◗
- FG (IEEE Int'l Conf. on Automatic Face and Gesture Recognition) ★
- IPDPS (IEEE Int'l Parallel and Distributed Processing Symposium) ◗

23 May

- ICSE (IEEE/ACM Int'l Conf. on Software Eng.) ▲

30 May

- ISCA (ACM/IEEE Int'l Symposium on Computer Architecture) ●

June

14 June

- CVPR (IEEE Conf. on Computer Vision and Pattern Analysis) ◗

16 June

- EuroS&P (IEEE European Symposium on Security & Privacy) ●

19 June

- JCDL (ACM/IEEE Joint Conf. on Digital Libraries) ▲

29 June

- DSN (IEEE/IFIP Int'l Conf. on Dependable Systems and Networks) ●

30 June

- MDM (IEEE Int'l Conf. on Mobile Data Management) ●

July

6 July

- ICME (IEEE Int'l Conf. on Multimedia and Expo) ●

8 July

- ICDCS (IEEE Int'l Conf. on Distributed Computing Systems) ▲

13 July

- COMPSAC (IEEE Annual Computer Software and Applications Conference) ●

August

31 August

- RE (IEEE Int'l Requirements Eng. Conf.) ●

September

21 September

- ASE (IEEE/ACM Int'l Conf. on Automated Software Eng.) ◆

28 September

- ICSME (IEEE Int'l Conf. on Software Maintenance and Evolution) ◆
- SecDev (IEEE Secure Development) ◗

October

18 October

- MODELS (ACM/IEEE Int'l Conf. on Model Driven Eng. Languages and Systems) ◗

21 October

- FIE (IEEE Frontiers in Education Conf.) ●

25 October

- VIS (IEEE Visualization Conf.) ◗

November

9 November

- FOCS (IEEE Annual Symposium on Foundations of Computer Science) ◗

15 November

- SC ◗

16 November

- LCN (2020 IEEE Conf. on Local Computer Networks) ◆

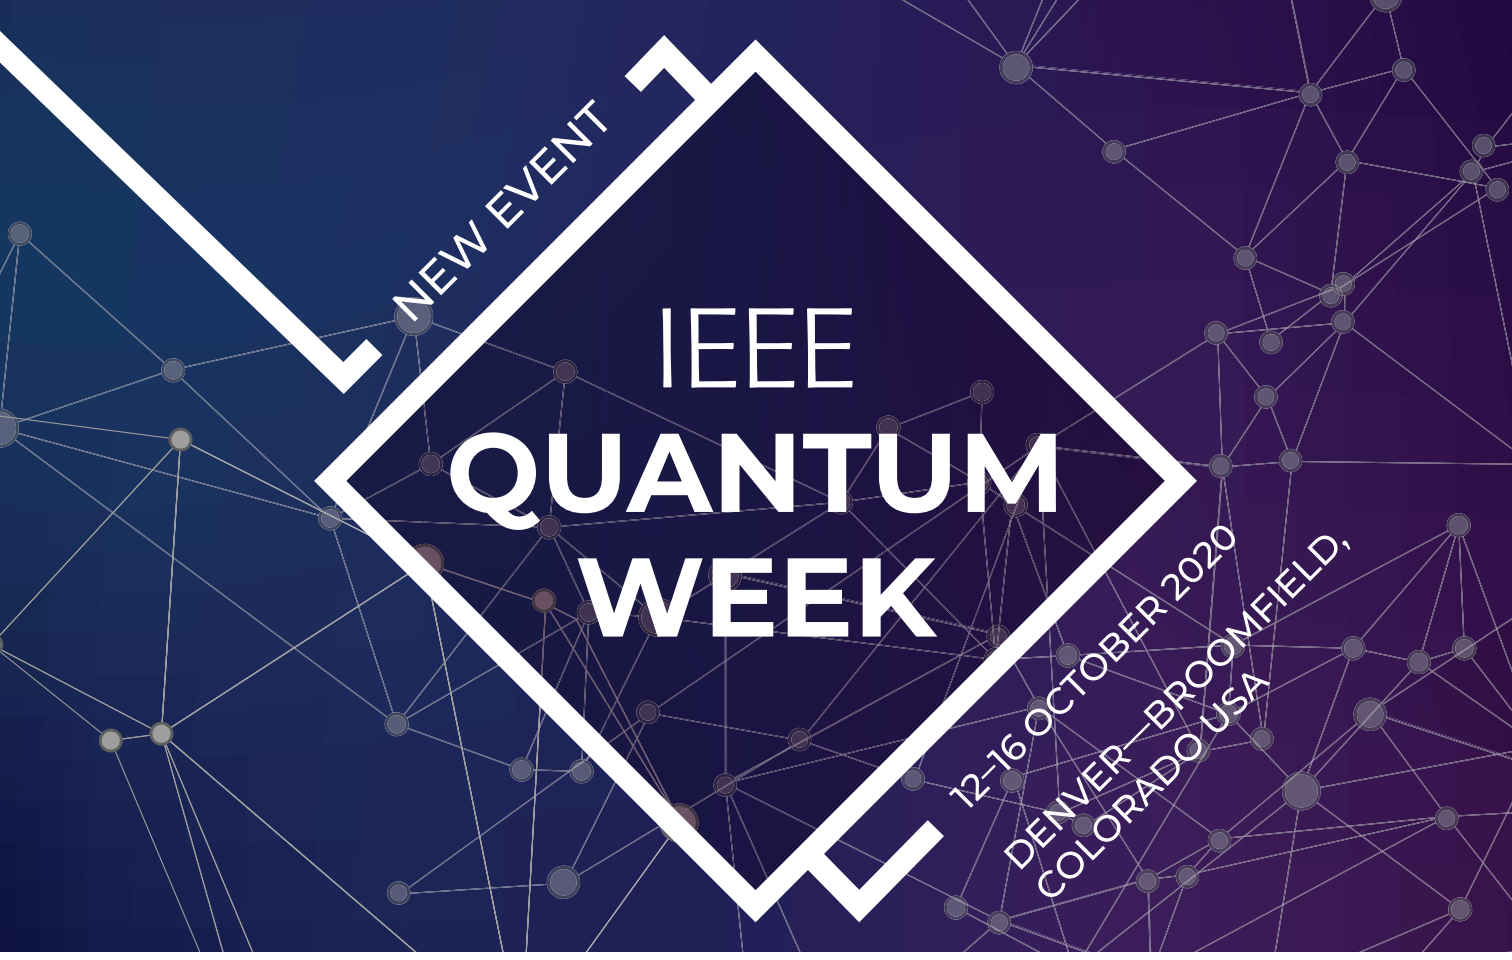

IEEE Quantum Week 2020 Is Open for Submissions

Participation opportunities are available for the inaugural IEEE International Conference on Quantum Computing and Engineering (QCE 2020) to be held 12–16 October 2020, in Denver—Broomfield, CO.

IEEE Quantum Week aims to be a leading venue for presenting high-quality original research, ground-breaking innovations, and compelling insights in quantum computing, engineering, and technologies.

Authors are invited to submit proposals for technical papers, posters, tutorials, workshops, and panels. Submission schedules are available at qce.quantum.ieee.org/important-dates.

IEEE Quantum Week includes the following technical paper tracks:

- Quantum Communications, Sensing, Cryptography
- Quantum Photonics and Optics
- Quantum Computing
- Quantum Algorithms & Information
- Quantum Applications and Simulating Nature
- Quantum Engineering
- Quantum Benchmarks & Measurements
- Quantum Education

Papers accepted by IEEE QCE will be submitted to the IEEE Xplore Digital Library. The best papers will be invited to the journals *IEEE Transactions on Quantum Engineering* (*TQE*) and *ACM Transactions on Quantum Computing* (*TQC*).

Submission instructions and details:
qce.quantum.ieee.org/callforcontributions

IEEE QUANTUM · IEEE COMPUTER SOCIETY · IEEE ComSoc IEEE Communications Society · IEEE CSC Council on Superconductivity · IEEE Photonics Society · IEEE