COMPUTING edge

ARTIFICIAL INTELLIGENCE

Also in this issue:

> Technology Policy and the Trump Administration

> The First Female Engineer at JPL

FEBRUARY 2018

www.computer.org

IEEE

IEEE computer society

IEEE computer society

SkillChoice™ Complete

Now with expanded libraries and an upgraded platform!

Valued at **$3,300!**

3,000+ online courses

MENTORSHIP

28,000+ BOOKS

6,000+ VIDEOS

Practice Exams

15,000+ *Books24x7 titles*

OVER 20x as many resources as before

One membership. **Unlimited knowledge.**

Did you know IEEE Computer Society membership comes with access to a high-quality, interactive suite of professional development resources, available 24/7?

Powered by Skillsoft, the SkillChoice™ Complete library contains more than $3,000 worth of industry-leading online courses, books, videos, mentoring tools and exam prep. Best of all, you get it for the one low price of your Preferred Plus, Training & Development, or Student membership package. There's something for everyone, from beginners to advanced IT professionals to business leaders and managers.

The IT industry is constantly evolving. Don't be left behind. Join the IEEE Computer Society today, and gain access to the tools you need to stay on top of the latest trends and standards.

Learn more at **www.computer.org/join**.

ACCESS TO SKILLSOFT IS AVAILABLE WITH

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
Sumi Helal, *Lancaster University*

IEEE Software
Diomidis Spinellis, *Athens University of Economics and Business*

IEEE Internet Computing
M. Brian Blake, *University of Miami*

IT Professional
Irena Bojanova, *NIST*

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

IEEE Micro
Lieven Eeckhout, *Ghent University*

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

IEEE Pervasive Computing
Marc Langheinrich, *University of Vienna*

Computing in Science & Engineering
Jim X. Chen, *George Mason University*

IEEE Intelligent Systems
V.S. Subrahmanian, *University of Maryland*

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

IEEE Annals of the History of Computing
Nathan Ensmenger, *Indiana University Bloomington*

IEEE Cloud Computing
Mazin Yousif, *T-Systems International*

FEBRUARY 2018 • VOLUME 4, NUMBER 2

COMPUTING edge

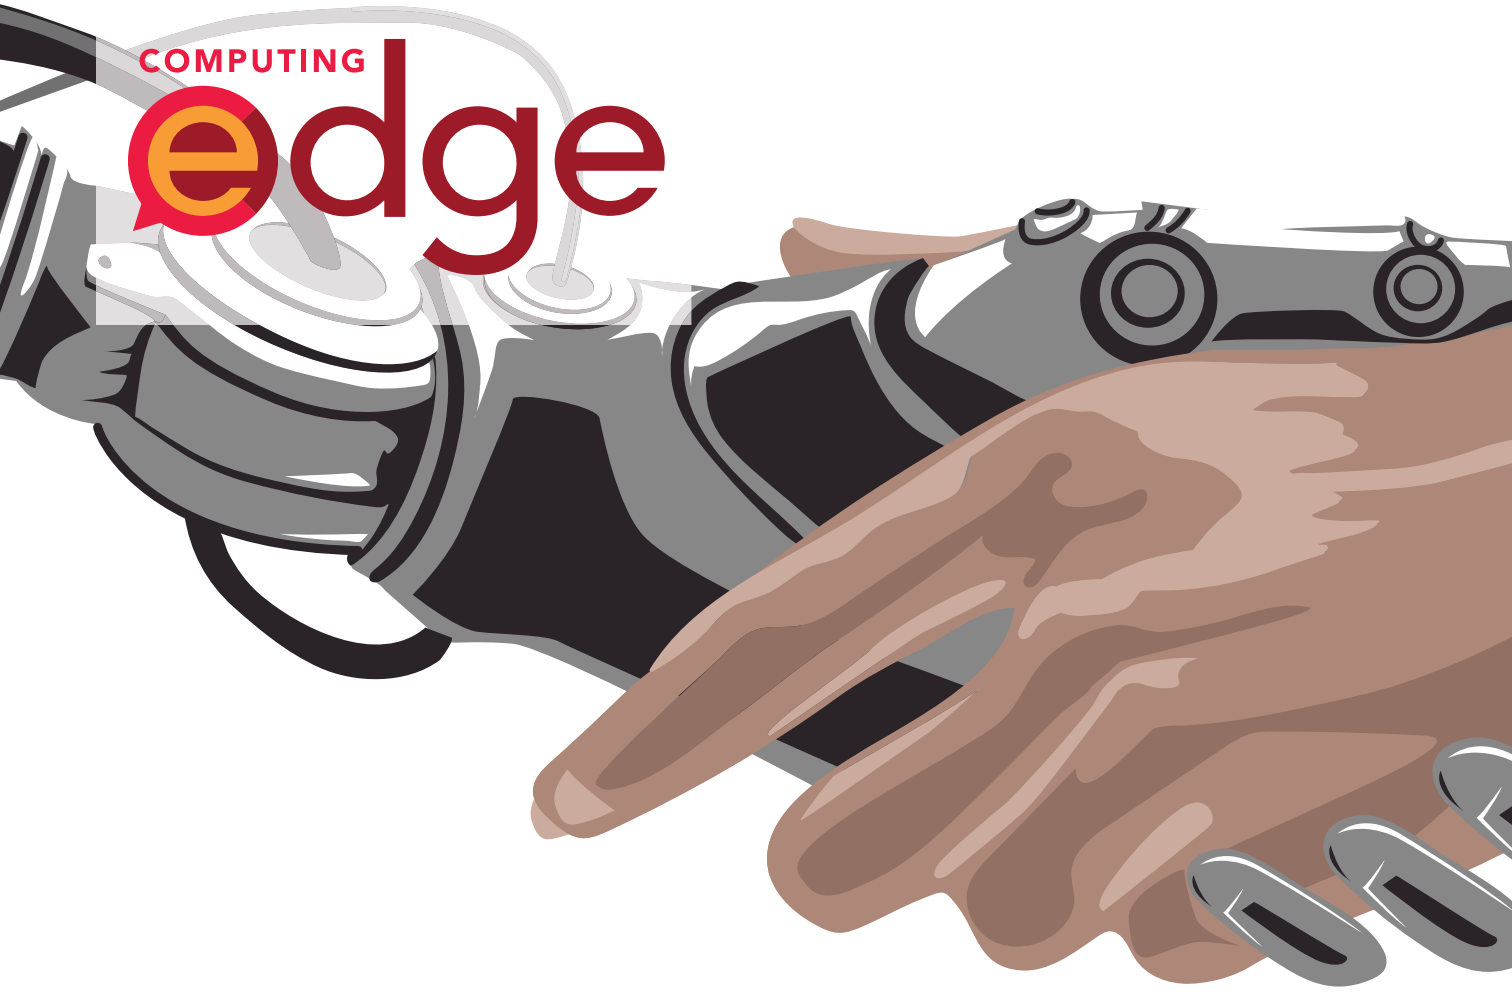

41

The Problem with AI

Departments

Magazine Roundup

by Lori Cameron

The IEEE Computer Society's lineup of 13 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 cloud migration and microchip design. Here are highlights from recent issues.

Computer

Current US Federal Policy Framework for Self-Driving Vehicles: Opportunities and Challenges

Self-driving systems raise many economic, social, and security issues that draw federal attention. The landscape is constantly shifting as US government agencies adapt policies, regulations, and funding priorities to address emerging technologies. Although many developers and technologists bemoan the complexity, pace, and ineptitude of government, others familiarize themselves with the challenges and seize opportunities in the policy arena. In this article from the December 2017 issue of *Computer*, researchers survey current federal policies and activities impacting technology developers, with special emphasis on privacy, cybersecurity, safety regulation, energy and environment, and ethical issues.

Computing in Science & Engineering

Product Innovation through Computational Prototypes and Supercomputing

Converting from physical prototype–based product design to computational (virtual) prototype–based product design requires leading-edge computational engineering codes; joint

Published by the IEEE Computer Society

R&D efforts; dedicated researchers; and meticulous verification, validation, and uncertainty quantification. For the most difficult engineering design problems, most of these are necessary but not sufficient. In extreme cases, an impending crisis might force conversion even in the face of strong resistance. This article from the November/December 2017 issue of *Computing in Science & Engineering* illustrates the Goodyear Tire & Rubber Company's successful, crisis-driven transition to virtual prototype–based product design.

IEEE Annals of the History of Computing

Carries Stripped to the Bone: Episodes in the History of Coaxial Modular Digital Counters

Although much has been written on the history of calculating machines, very little attention has been paid to the evolution of mechanical counters and their components. Mechanical counters were ubiquitous and could be found in many cars, where they served as odometer displays, or in cash registers and various calculating machines. The most common construction for such counters is made of rotating disks, which are similar and located on the same axis. Although they look simple in appearance and perhaps standardized, they have a history of their own. In this article from the July–September issue of *IEEE Annals of the History of Computing*, the author analyzes three of the earliest known models of counters, which can be viewed as ancestors of the modern mechanical counter.

IEEE Cloud Computing

Realizing Software Reliability in the Face of Infrastructure Instability

Cloud computing has brought with it the utilization of off-the-shelf commodity hardware that has higher failure rates than the systems used in enterprises for the past several decades. Coupled with increasingly complex, highly distributed, constantly changing datacenter environments that can no longer be treated as deterministic systems, this forces us to change the way we view the stability of that infrastructure. This article from the September/October 2017 issue of *IEEE Cloud Computing* studies how Netflix has fully embraced this mindset change and was able to avoid any significant impact during a major outage experienced by their cloud infrastructure provider, Amazon Web Services (AWS).

Experiencing the Sights, Smells, Sounds, and Climate of Southern Italy in VR

In this article from the November/December 2017 issue of *IEEE Computer Graphics and Applications*, researchers explore what it takes to make interactive computer graphics and virtual reality (VR) attractive as a promotional vehicle, from the points of view of the tourism agencies and the tourists themselves. Specifically, in response to a call from local authorities seeking to increase the tourism appeal of the Apulia region in southern Italy, the authors proposed an alternative approach to traditional tourism marketing and advertising efforts: an interactive, innovative, and attractive VR experience called the Multisensory Apulia Touristic Experience (MATE).

IEEE Intelligent Systems

Design and Prototyping a Smart Deep Brain Stimulator: An Autonomous Neuro-Sensing and Stimulating Electrode System

In this article from the September/October 2017 issue of *IEEE Intelligent Systems*, researchers present the design and prototyping of a smart deep brain stimulator (SDBS) that consists of brain-implantable smart electrodes and a wireless-connected external controller. SDBS electrodes operate as completely autonomous electronic implants that are capable of sensing and recording neural activities in real time, performing local processing, and generating arbitrary waveforms for neuro-stimulation. A bidirectional, secure, fully passive wireless communication backbone was designed and integrated into this smart electrode to maintain contact between the electrodes and the controller.

IEEE Internet Computing

PACMAN: Personal Agent for Access Control in Social Media

Given social media users' plethora of interactions, appropriately controlling who can access what information becomes a challenging task

for users. Selecting the appropriate audience, even from within their own friend network, can be fraught with difficulties. In this article from the November/December 2017 issue of *IEEE Internet Computing*, researchers present PACMAN as a potential solution. It's a personal assistant agent that recommends personalized access control decisions based on the social context of any information disclosure by incorporating communities generated from the user's network structure and utilizing information in the user's profile. PACMAN provides accurate recommendations while minimizing intrusiveness.

IEEE Micro

Flying IoT: Toward Low-Power Vision in the Sky

Many Internet of Things (IoT) devices require some level of machine learning or cognitive capability to be truly effective, but the high computational complexity of cognitive algorithms makes them unsuitable for low-power IoT processors. In this article from the November/December 2017 issue of *IEEE Micro*, the authors study a cognitive drone application in a design space called Flying IoT to better understand the design challenges of cognitive IoT devices. To improve their processor's performance while maintaining its low-power consumption and small form factor, the authors propose a sensor-cloud architecture in which data collection is done at the edge and data processing is offloaded to the cloud. This architecture can process complex convolution neural network models in near real time with software optimizations such as downsampling and compression, while consuming less power than state-of-the art embedded processors such as the Jetson TX1.

IEEE MultiMedia

Crowdsensing Multimedia Data: Security and Privacy Issues

Smartphones are equipped with various sensors and have high-performance wireless communication capabilities. Through the ubiquitous presence of powerful mobile devices, crowdsensing lets ordinary people collectively gather and share real-time multimedia data. Multimedia crowdsensing has made large-scale participatory sensing viable in a speedy and cost-efficient manner, but it also introduces security and privacy concerns. For example, participants' personally identifiable information can be exposed while sharing individually owned sensor data. In this article from the October–December 2017 issue of *IEEE MultiMedia*, the authors identify security and privacy issues in multimedia crowdsensing and describe existing solutions that are designed to protect both data producers and consumers.

IEEE Pervasive Computing

Typhlex: Exploring Deformable Input for Blind Users Controlling a Mobile Screen Reader

Current smartphone technology presents many challenges for blind users. The authors of this article from the October–December 2017 issue of *IEEE Pervasive Computing* introduce the use of a deformable device prototype, Typhlex, with strategically placed grooves to elicit bend gestures. They conducted two exploratory studies with sighted participants (with the prototype hidden from view) and blind participants, focusing on comparing the usability of bend gestures to touch as primary forms of input. Their findings suggest that while easily learnable and enjoyed by both groups, the prototype had yet to improve blind users' performances when compared to the commonly used touch input paradigm. They present lessons learned from their design process and studies, and discuss the promise of deformable input devices in the area of accessibility for blind users.

IEEE Security & Privacy

Botnet Fingerprinting: Anomaly Detection in SMTP Conversations

In this article from the November/December 2017 issue of *IEEE Security & Privacy*, the authors present the results obtained during their research on detecting unsolicited emails sent by botnets. The distinction from most existing solutions is that this approach is based on the analysis of network traffic, specifically the sequence and syntax of SMTP commands observed during email delivery. The authors present several improvements for detecting unsolicited email sources from different

botnets (fingerprinting) that can be used during network forensic investigation.

IEEE Software

Probabilistic Threat Detection for Risk Management in Cyber-Physical Medical Systems

Medical devices are complex cyber-physical systems incorporating emergent hardware and software components. However, this complexity leads to a wide attack surface, posing security risks and vulnerabilities. Mitigation and management of such risks during premarket design and postmarket deployment are required. In this article from the January/February 2018 issue of *IEEE Software*, the authors present a dynamic risk management and mitigation approach based on probabilistic threat estimation. A smart pacemaker case study illustrates the approach.

IT Professional

Graph Databases for Knowledge Management

Emerging technologies let companies manage their knowledge assets with more innovative and effective methods. Due to the complex nature of knowledge management processes, it is cumbersome to design, develop, and implement a system based on relational databases. This article, which appears in the November/December 2017 issue of *IT Pro*, proposes a specific graph database application in streamlining major knowledge management processes. The author develops a property graph data model to facilitate the process model of knowledge management. In addition, the model is implemented through the Neo4j graph database system. This research provides some guidance for practitioners in seeking alternative approaches to traditional methods of knowledge management.

Computing Now

The Computing Now website (computingnow.computer.org) features up-to-the-minute computing news and blogs, along with articles ranging from peer-reviewed research to opinion pieces by industry leaders.

Artificial Intelligence

From Siri and Alexa to self-driving cars and robots fulfilling Amazon orders, artificial intelligence already surrounds us. As AI advances and permeates more areas of life, there are obvious benefits and drawbacks—superintelligence could help us eradicate war, poverty, and disease, but it could also intentionally or unintentionally cause great harm. The February 2018 issue of *ComputingEdge* features cutting-edge research and thought pieces on AI and more.

In *Computer*'s "Standardizing Ethical Design for Artificial Intelligence," the authors recognize that AI's consequences for social order aren't well understood, and look at how standards can guide the way technology impacts society.

Researchers at Harvard University propose a conceptual framework for thinking about governance for AI in *IEEE Internet Computing*'s "A Layered Model for AI Governance," noting that there is a large information gap between AI developers and consumers and policymakers.

In *IEEE MultiMedia*'s "Deep Learning Triggers a New Era in Industrial Robotics," the author examines deep learning applications in robotics and introduces promising research that will likely impact how industrial robot systems are designed in the near future.

The author of *IT Professional*'s "How Much to Trust Artificial Intelligence?" notes that AI is typically software dominant, and that software is prone to vulnerabilities. He recommends using caution until some reliable methodology is adopted for the assessment of assured trust within AI.

In *IEEE Pervasive Computing*'s "Augmenting Human Intellect and Amplifying Perception and Cognition," the author examines various technologies designed to augment human intellect and amplify human perception and cognition, considering how novel technologies can create a new relationship between digital technologies and humans.

In *IEEE Software*'s "Katie Malone on Machine Learning," Edaena Salinas interviews Katie Malone, a data scientist in the R&D department at Civis Analytics, which specializes in data science software and consulting. Katie and Edaena discuss the major types of machine-learning algorithms—the backbone of AI—and some examples, including supervised and unsupervised classification.

Finally, in *IT Professional*'s "The Problem with AI," the author looks at how AI and machine-learning technologies can help or hinder organizations in curating vast amounts of data.

This *ComputingEdge* issue also includes articles on topics other than AI:

- The author of *IEEE Micro*'s "Technology Policy and the Trump Administration" considers how Trump's (likely) technology policies will affect the value of US firms in IT markets.
- *IEEE Cloud Computing's* "Cloud-Native Applications and Cloud Migration: The Good, the Bad, and the Points Between" looks at the advantages and costs of cloud-native features in IT.
- *IEEE Annals of the History of Computing*'s "Dana Ulery: Pioneer of Statistical Computing and Architect of Large, Complex Systems" profiles JPL's first female engineer.
- *CiSE*'s "Computing in World War I" looks back at the Royal Navy's use of computers in their battleships.

 Published by the IEEE Computer Society 2469-7087/18/$33.00

IEEE COMPUTER SOCIETY: Be at the Center of It All

IEEE Computer Society membership puts you at the heart of the technology profession—and helps you grow with it.

Here are 10 reasons why you need to belong.

IEEE Computer Society—keeping you ahead of the game. Get involved today.

www.computer.org/membership

IEEE computer society

Standardizing Ethical Design for Artificial Intelligence and Autonomous Systems

Joanna Bryson, University of Bath

Alan Winfield, University of the West of England

AI is here now, available to anyone with access to digital technology and the Internet. But its consequences for our social order aren't well understood. How can we guide the way technology impacts society?

For decades—even prior to its inception—AI has aroused both fear and excitement as humanity has contemplated creating machines like ourselves. Unfortunately, the misconception that "intelligent" artifacts should necessarily be human-like has largely blinded society to the fact that we have been achieving AI for some time. Although AI that surpasses human ability grabs headlines (think of Watson, Deep Mind, or alphaGo), AI has been a standard part of the industrial repertoire since at least the 1980s, with expert systems checking circuit boards and credit card transactions.

Machine learning (ML) strategies for generating AI have also long been used, such as genetic algorithms for finding solutions to intractable computational problems like scheduling, and neural networks not only to model and understand human learning but also for basic industrial control, monitoring, and classification. In the 1990s, probabilistic and Bayesian methods revolutionized ML and opened the door to one of the most pervasive AI abilities now available: searching through massive troves of data. Innovations in AI and ML algorithms have extended our capacity to find information in texts, allowing us to search photographs as well as both recorded and live video and audio. We can translate, transcribe, read lips, read emotions (including lying), forge signatures and other handwriting, and forge video.

Yet, the downside of these benefits is ever present. As we write this, allegations are circulating that the

 Published by the IEEE Computer Society 2469-7087/18/$33.00

EDITOR **CHUCK WALRAD**
Davenport Consulting; cwalrad@daven.org

outcomes of the recent US presidential election and UK referendum on EU membership were both influenced by the use of AI to detect and target "swing voters" via public social media. To address these and other concerns, the IEEE Computer Society Standards Activities Board is creating standards for responsible designers who will shape our brave new world and ensure AI's benefit to humanity.

DEFINING AI

Although the following definitions are not universally used, they're well-established.[1] *Intelligence* is the capacity to do the right thing at the right time, in a context where doing nothing (or making no change in behavior) would be worse. Intelligence then requires

- the capacity to perceive contexts for action,
- the capacity to act, and
- the capacity to associate contexts to actions.

By this definition, plants are intelligent. They can perceive and respond to the direction of light, for example. The more conventional understanding of "intelligent" includes being cognitive, that is, being able to learn new contexts and actions, and the associations between them.

AI, by convention, describes (typically digital) artifacts that demonstrate any of these capacities. So, for example, machine vision, speech recognition, pattern recognition, and static production systems are all examples of AI, with algorithms that can be found in standard AI textbooks.

Robots are artifacts that sense and act in the physical world in real time. By this definition, a smartphone is a (domestic) robot. It has not only microphones but also a variety of proprioceptive sensors that let it know when its orientation is changing or when it is falling.

Autonomy is technically the capacity to act as an individual. For social animals like humans, autonomy is normally situated somewhere along a scale. For example, it is fully expected that family, workplace, government, and other organizations might regularly have some impact on our actions. Similarly, a technical system that can sense the world and select an action specific to its present context is called "autonomous" even though its actions are ultimately determined by the designers that constructed its intelligence and its operators.

CONCERNS ABOUT DOMESTIC AND COMMERCIAL AI

AI is core to some of the most successful companies in history in terms of market capitalization and, along with information and communications technology (ICT) more generally, has revolutionized the ease with which people from all over the world can create, access, and share knowledge. However, possible pitfalls of AI could have quite serious consequences. Here we briefly review some common concerns to see which are both realistic and specific to AI.

Will AI outcompete us?

Some of the most sensational fears are that, as AI increases to the point that it surpasses human abilities, it might take control over our resources and outcompete our species, leading to human extinction. AI is already superhuman in many domains. With machines, we can already do arithmetic better, play chess and Go better, transcribe speech better, read lips better, remember more things for longer, and indeed be faster and stronger than we are unaided. However, these capacities have in no sense led to machine ambition. Human memory has been outstripped by books for centuries—mere intelligence is no more of a direct threat than mere strength.

Will AI undermine societal stability?

For centuries, people have had significant concerns about the displacement of workers by technology. There is no question that new technologies disrupt communities, families, and lives, but historically, the majority of this disruption has been for the better. In general, lifespans are longer and infant mortality is lower than ever before, and these indicators are well associated with political stability. Nevertheless, we are currently seeing a disruption that seems to be undermining political stability. This disturbance is termed *political polarization*, which seems to co-occur with inequality, although causality between these is unclear.[2] Polarization has happened before, for example, in the early 20th century, reaching its climax in World War I. New technologies could play a role in increasing inequality—and therefore polarization—by eliminating costs such as distance that formerly supported economic diversity. This time, AI and ICT might be the technologies changing the economic landscape.

Will AI harm privacy, personal liberty, and autonomy?

What really makes AI special is its relationship to information, especially personal information. Previous periods of domestic spying have been associated with everything from prejudice in opportunities to pogroms. However, AI and ICT can greatly facilitate such knowledge gathering. We are now able to keep and access long-term records on anyone who produces storable data—for example, anyone with bills, contracts, or a credit history, not to mention public writing and social

media use. With ML, this data lets us make predictions concerning individuals' behavior and preferences, which in turn opens the possibilities of control or persecution.

CAN STANDARDS PROMOTE ETHICS IN AI?

Standards are consensus-based agreed-upon ways of doing things, setting out how things should be done. If a system or process can be shown to do things as prescribed, it is said to be compliant with the standard. Such compliance provides confidence in a system's efficacy in areas important to users, such as safety, security, and reliability.

> IEEE's Initiative for Ethical Considerations in Artificial Intelligence Systems has as its mission to "ensure every technologist is educated, trained, and empowered to prioritize ethical considerations."

Few standards explicitly address ethics in robotics and AI. One that does is British Standard (BS) 8611:2016, *Robots and Robotic Devices: Guide to the Ethical Design and Application of Robots and Robotic Systems*.[3] Published in April 2016, it provides designers with a tool to assess ethical risk. At the heart of BS 8611:2016 is a set of 20 distinct ethical hazards and risks, grouped under four categories: societal, application, commercial/financial, and environmental. Advice on measures to mitigate the impact of each risk is given, along with suggestions on how such measures might be verified or validated.

IEEE's Initiative for Ethical Considerations in Artificial Intelligence and Autonomous Systems, a program designed to bring together "multiple voices in the AI and Autonomous Systems (AS) communities," has as its mission to "ensure every technologist is educated, trained, and empowered to prioritize ethical considerations in the design and development of autonomous and intelligent systems."[4] The first output from the initiative is a discussion document called *Ethically Aligned Design* (EAD), version 1, published in December 2016.[4] The work of eight committees, it covers

- general principles,
- how to embed values into autonomous intelligent systems,
- methods to guide ethical design and design,
- safety and beneficence of artificial general intelligence and artificial superintelligence,
- personal data and individual access control,
- how to reframe autonomous weapons systems,
- economics and humanitarian issues, and
- law.

EAD articulates a set of about 60 draft issues and recommendations. Each committee was asked to identify issues that could be addressed through a new standard. Presently, four standards working groups are drafting candidate standards to address an ethical concern articulated by one or more of the eight committees outlined in the EAD document. The candidate standards are

- *P7000—Model Process for Addressing Ethical Concerns during System Design* (standards.ieee.org/develop/project/7000.html), which aims to establish a value-based system design methodology;
- *P7001—Transparency of Autonomous Systems* (standards.ieee.org/develop/project/7001.html), which we discuss below;
- *P7002—Data Privacy Process* (standards.ieee.org/develop/project/7002.html), which aims to create one overall methodological approach that specifies practices to manage privacy issues; and
- *P7003—Algorithmic Bias Considerations* (standards.ieee.org/develop/project/7003.html), which aims to specify methodologies to ensure that negative bias in algorithms is addressed and eliminated.

CASE STUDY: A STANDARD FOR TRANSPARENCY

P7001 is an effort in which both authors are involved. It is based on the radical proposition that it should always be possible to find out why an AS made a particular decision.

Transparency is not one thing. Clearly, elderly persons don't require the same level of understanding of their care robot as the engineer who repairs it. Nor would patients expect the same appreciation of the reasons a medical-diagnosis AI recommends a particular course of treatment as their doctor. The P7001 working group has identified five categories of stakeholder—users, safety certification agencies, accident investigators, lawyers or expert witnesses, and wider society—and proposes that ASs must be transparent to each in different ways and for different reasons.

- For users, transparency is important because it builds trust in the system by providing a simple way for users to understand what the system is doing and why.
- For AS safety certification, transparency is important because it exposes the system's processes for independent certification against safety standards.
- If accidents occur, an AS needs to be transparent to investigators; the internal process that

This article originally appeared in Computer, *vol. 50, no. 5, 2017.*

led to the accident must be traceable.
- Following an accident, lawyers or other expert witnesses who might be called on to give evidence require transparency to inform their evidence.
- Disruptive technologies, such as driverless cars, require a certain level of transparency for wider society to gain the public's confidence in the technology and to ensure that trust is deserved.

Of course, the way in which transparency is provided is likely to be very different for each group. If we take a care robot as an example, transparency means users can understand what the robot might do in different circumstances. If the robot does anything unexpected, they should be able to ask it "Why did you just do that?" and receive an intelligible reply.

Safety certification agencies will need access to technical details of how the AS works, together with verified test results. Accident investigators will need access to data logs of exactly what happened prior to and during an accident, most likely provided by something akin to an aircraft flight data recorder—and it should be illegal to operate an AS without such a system. Wider society would need accessible documentary-type science communication to explain an AS (such as a driverless car autopilot) and how it works.

In P7001, we aim to develop a standard that sets out measurable, testable levels of transparency in each of these categories (and perhaps new categories yet to be determined) so that we can assess an AS objectively and determine compliance. It is our aim that P7001 will also articulate transparency levels in a range that defines minimum levels up to the highest achievable standards of acceptance. The standard will provide AS designers with a toolkit for self-assessing transparency as well as recommendations for how to address shortcomings or transparency hazards.

The changes artificial intelligence and autonomous systems are bringing to the world are real, and already in progress. Although we cannot say with certainty that the situation is in hand, we as members of the global initiative are optimistic that the right steps are being taken and that IEEE will be key to ensuring that AI and ASs benefit all of humanity. C

ACKNOWLEDGMENTS

We thank John C. Havens for introducing us to the IEEE initiative and forwarding Dr. Walrad's request for an article to us.

REFERENCES

1. P.H. Winston, *Artificial Intelligence*, Addison-Wesley, 1984.
2. N.M. McCarty, K.T. Poole, and H. Rosenthal, *Polarized America: The Dance of Ideology and Unequal Riches*, MIT Press, 2006.
3. *Robots and Robotic Devices: Guide to the Ethical Design and Application of Robots and Robotic Systems*, BS 8611:2016, British Standards Inst., 2016; shop.bsigroup.com/ProductDetail?pid=000000000030320089.
4. *Ethically Aligned Design: A Vision for Prioritizing Human Wellbeing with Artificial Intelligence and Autonomous Systems*, version 1, IEEE Standards Assoc., 2016; standards.ieee.org/develop/indconn/ec/ead_v1.pdf.

JOANNA BRYSON is an associate professor in the Department of Computer Science at the University of Bath and an affiliate of the Princeton Center of Technology Policy. Contact her at jjb@alum.mit.edu.

ALAN WINFIELD is a professor of robot ethics at the Bristol Robotics Laboratory in the University of the West of England and a visiting professor in the Department of Electronic Engineering at the University of York. Contact him at alan.winfield@uwe.ac.uk.

Editor: Virgilio A.F. Almeida • virgilio@dcc.ufmg.br

A Layered Model for AI Governance

Urs Gasser and Virgilio A.F. Almeida • *Harvard University*

AI-based systems are "black boxes," resulting in massive information asymmetries between the developers of such systems and consumers and policymakers. In order to bridge this information gap, this article proposes a conceptual framework for thinking about governance for AI.

Many sectors of society rapidly adopt digital technologies and big data, resulting in the quiet and often seamless integration of AI, autonomous systems, and algorithmic decision-making into billions of human lives.[1,2] AI and algorithmic systems already guide a vast array of decisions in both private and public sectors. For example, private global platforms, such as Google and Facebook, use AI-based filtering algorithms to control access to information. AI algorithms that control self-driving cars must decide on how to weigh the safety of passengers and pedestrians.[3] Various applications, including security and safety decision-making systems, rely heavily on AI-based face recognition algorithms. And a recent study from Stanford University describes an AI algorithm that can deduce the sexuality of people on a dating site with up to 91 percent accuracy.[4] Voicing alarm at the capabilities of AI evidenced within this study, and as AI technologies move toward broader adoption, some voices in society have expressed concern about the unintended consequences and potential downsides of widespread use of these technologies.

To ensure transparency, accountability, and explainability for the AI ecosystem, our governments, civil society, the private sector, and academia must be at the table to discuss governance mechanisms that minimize the risks and possible downsides of AI and autonomous systems while harnessing the full potential of this technology.[5] Yet the process of designing a governance ecosystem for AI, autonomous systems, and algorithms is complex for several reasons. As researchers at the University of Oxford point out,[3] separate regulation solutions for decision-making algorithms, AI, and robotics could misinterpret legal and ethical challenges as unrelated, which is no longer accurate in today's systems. Algorithms, hardware, software, and data are always part of AI and autonomous systems. To regulate ahead of time is difficult for any kind of industry. Although AI technologies are evolving rapidly, they are still in the development stages. A global AI governance system must be flexible enough to accommodate cultural differences and bridge gaps across different national legal systems. While there are many approaches we can take to design a governance structure for AI, one option is to take inspiration from the development and evolution of governance structures that act on the Internet environment. Thus, here we discuss different issues associated with governance of AI systems, and introduce a conceptual framework for thinking about governance for AI, autonomous systems, and algorithmic decision-making processes.

The Nature of AI

Although AI-based applications are increasingly adopted in hospitals, courtrooms, schools, at home, and on the road to support (and in some instances, even guide) human decision-making,

 Published by the IEEE Computer Society

currently there is no universally accepted definition of AI, a term coined in the mid-1950s by US researchers.[6,7] One reason for the lack of a definition is that AI, from a technical perspective, is not a single technology, but rather a set of techniques and subdisciplines ranging from areas such as speech recognition and computer vision to attention and memory, to name just a few.[6]

From a phenomenological perspective, however, the term AI is often used as an umbrella term to refer to a certain degree of autonomy exhibited in advanced health diagnostic systems, next-generation digital tutors, self-driving cars, and other AI-based applications. Often, such applications in turn impact human behavior and evolve dynamically in ways that are at times unforeseen by the systems' designers. In this context, the differentiation between weak (or narrow) and strong (or general) AI is often used and helpful when discussing the nature of AI. Weak AI describes the current generation of applications that are focused on a relatively narrow task such as playing a game, recognizing a voice, or detecting certain patterns on a CT-scan. Strong AI, in contrast, refers to machines with genuine intelligence and self-awareness in the sense that the machine has the ability to apply intelligence to any problem.[8] At present, the technical possibility and (potential) societal impact of strong AI is discussed controversially, while the current adaptation of weak AI already leads to a series of real governance issues that deserve attention in the present.

AI Governance Challenges

Following a typical pattern when new technologies become more widely available, policymakers and other stakeholders are focusing largely on the risks and harms of AI-based technologies.[9] Again, similar to previous conversations about digital technologies' impact on society, the challenges related to AI, autonomous systems, and algorithms are often presented and discussed in the form of lists of substantive issues (including policy, legal, governance, and ethical considerations) that must be addressed.[6]

A recent roadmap on AI policy by one leading expert, for instance, identifies the following clusters of core issues and questions where AI applications either lead to new challenges or amplify pre-existing policy concerns and pressure points[10]:

- *Justice and equality*. To what extent can AI systems be designed and operated to reflect human values such as fairness, accountability, and transparency and avoid (new) inequalities and biases?
- *Use of force*. As AI-based systems are now involved in making decisions about the use of force — for instance, in the case of autonomous weapons — how much human control is necessary or required? Who bears responsibility for the AI-based outputs?
- *Safety and certification*. Particularly where AI-based systems have a physical manifestation, how do we define and validate safety thresholds — for instance, through standard-setting and certification?
- *Privacy*. As AI-systems are enabled and powered by data, what are the privacy implications and new privacy threats of next-generation technologies — for instance, in terms of government surveillance or corporate influence over customers?
- *Displacement of labor and taxation*. To what extent will AI-based machines replace jobs previously performed by humans, or at least transform what labor means? What are the effects of AI on public finances if robots don't pay taxes?

Such lists of substantive issues, to which several others could be added (for instance, intellectual property or liability), can be supplemented by cross-cutting themes surrounding transparency, accountability, and explainability; inclusion and fairness; global governance; and more that span across the different application areas of AI-based systems (see https://cyber.harvard.edu/research/ai/usecases).

Models for AI Governance

When considering future governance models for AI that address the aforementioned issues, it might be helpful and necessary to move beyond such lists and consider some of the larger structural challenges associated with the "regulation" (broadly defined) of AI-based technologies. In the following, we highlight three such challenges that translate into design requirements for a future governance model of AI.

Information asymmetries. While AI has the potential to shape the lives of billions of people, only a few experts really understand the underlying techniques. AI-based systems are often inscrutable, sometimes resulting in massive information asymmetries between the developers of such systems and other stakeholders, including consumers and policymakers. An effective governance system for AI needs to incorporate mechanisms aimed at improving our collective understanding of the AI phenomenon in its different manifestations and contexts of application.

Finding normative consensus. The current policy and governance debate is largely focused on risks and challenges associated with AI. But AI also offers tremendous potential benefits to society, as the discussions

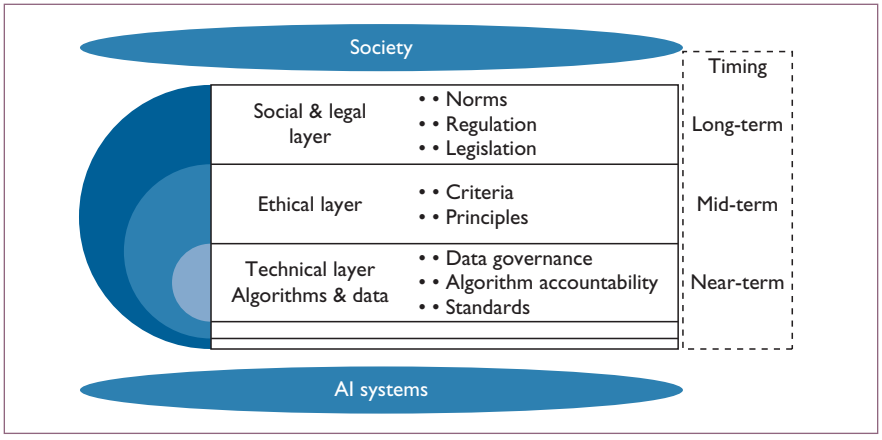

Figure 1. A layered model for AI governance. The interacting layers (which sit between society and AI applications) are social and legal; ethical; and technical foundations that support the ethical and social layers.

about the use of AI in the context of Sustainable Development Goals illustrate (see www.itu.int/en/ITU-T/AI/Pages/201706-default.aspx). A governance model must open up spaces for cost-benefit analyses and normative consensus building among different stakeholders, particularly where tradeoffs are involved in the design of AI systems. A future governance model also needs to deal with normative differences among contexts and geographies, and provide for interoperability among different frameworks and approaches.[11]

Government mismatches. Even where we have a shared understanding of AI technologies, the underlying techniques, and societal consensus about what is or isn't desirable, the design of effective, efficient, and legitimate means (strategies, approaches, tools, and so forth) to resolve the aforementioned substantive issues is challenging, given the conditions of uncertainty and complexity in the AI ecosystem. But larger undercurrents also put limits on traditional approaches to law- and policymaking in the digital age.[12]

Taken together, these structural challenges and associated design requirements for a future governance model of AI point away from simple state-centric, command-and-control regulatory schemes toward more complex approaches to governance emergent in fields as diverse as the Internet, nanotechnology governance, or gene driver governance. While the exact contours of a future AI governance model are still in flux, advanced governance models such as active matrix theory, polycentric governance, hybrid regulation, and mesh regulation can provide both inspiration and conceptual guidance on how such a future governance regime might be designed.[13] In the next section, we highlight one feature that is common across many of these models: the idea of modularity embodied in the form of layered governance, which also combines different instruments to grapple with and address the aforementioned substantive issues, making it a shared responsibility among all relevant actors involved. It is important to note that any such emerging model must be situated in and interact with existing institutional frameworks of applicable laws and policies, particularly human rights, as the development and deployment of AI does not take place in a vacuum.[14]

The Layered Model

Modularity is one of the main mechanisms for managing complex systems. Modularity aims to reduce the number of interdependencies that must be analyzed by identifying which tasks are highly interdependent and which ones are not.[15] Layering represents a particular form of modularity, in which different parts of the overall system are arranged into parallel hierarchies. A frequently cited example of layering is the Open System Interconnection (OSI) Reference model used during the late 1970s.[15] Another example of a layered model was proposed by David Clark[16] to represent the nature of cyberspace using a model with four layers, that are: first, the people who participate in the cyber-experience; second, the information that is stored, transmitted, and transformed in cyberspace; third, the logical building blocks that make up the services, and fourth, the physical foundations that support the logical elements. The scale, heterogeneity, complexity, and degree of technological autonomy of AI systems require new thinking about policy, law, and regulation. We attempt to capture the complex nature of AI governance by using an analytical model with three layers. From the top down, the interacting layers are as follows:

- social and legal;
- ethical; and
- technical foundations that support the ethical and social layers.

Figure 1 shows a representation of the layered governance model. It will sit between society and AI applications. The instruments mapped onto the layers can be developed at different times. In the near term, governance proposals could concentrate on developing standards and principles for AI algorithms. For the mid- and long-term, nation-states can work on specific legislation to

regulate mature AI applications. The model can be a helpful heuristic that illustrates how principles, policies, norms, and laws in response to AI-based challenges and opportunities can be combined and work together, within and across layers.

The Technical Layer

The technical layer is the foundation of the AI governance ecosystem — the algorithms and data out of which it is built. AI systems and autonomous systems rely on data and algorithms, regardless of whether they are physical systems (such as self-driving cars and commercial robots) or software systems (such as criminal justice or medical diagnostic systems, or intelligent personal assistants).[17] A set of principles for accountable algorithms and an associated suggested social impact statement were developed as part of a Dagstuhl Seminar on "Data, Responsibly."[18] The proposed principles for accountable algorithms with social impact are as follows: responsibility, explainability, accuracy, auditability, and fairness. The collection, use, and management of data by AI algorithms, known as data governance, should follow principles that promote fairness and safeguard against race, color, national origin, religion, sex, gender, sexual orientation, disability, or family status discrimination.[19]

The Ethical Layer

On top of the technical layer, we could articulate high-level ethical concerns that apply to different types of AI applications and systems. One important source for the development of such ethical principles are human rights principles. Another example of the emergence of AI ethics norms is the IEEE general principles for AI and autonomous systems.[17] Actions driven by algorithms can be assessed according to ethical criteria and principles. For instance, when an AI application analyzes the data of an insurance company and charges a certain group of people higher premiums, based on variables such as gender or age, such a decision-making application would be violating the ethical principle of equal or fair treatment.

The Social and Legal Layer

The social and legal layer could address the process of creating institutions and allocating responsibilities for regulating AI and autonomous systems. For example, Matthew Scherer[20] describes a policymaking body that would have the power to define AI, create exceptions allowing for AI research to be conducted in certain environments without the researchers being subjected to strict liability, and establish an AI certification process. One starting point for specific norms aimed at regulating AI can be the principles and criteria that emerge from the ethical and technical layers, in addition to pre-existing and more general national and international legal frameworks, including human rights. The layered model provides a framework for thinking about AI governance, aiming at the definition of appropriate behavior for AI and autonomous systems.

Implementing governance structures for AI and algorithmic decision-making systems can occur at multiple layers and involve blended approaches. Here, we describe some of these layers, taking into consideration that some of them would only be considered if the risks that certain AI applications present are substantial and concrete. Governance processes can range from market-oriented solutions to government-based structures and can be applied nationally or internationally. On the regional level, a rich example is the General Data Protection Regulation (GDPR), a wide-ranging and complex regulation intended to strengthen and unify data protection for all individuals within the European Union (www.eugdpr.org). It offers a (limited) "right to explanation" that will oblige companies to explain the purpose of an algorithm and the kind of data it uses when making automated decisions.[21] Absent an AI-specific international legal framework, a global oversight body, which can take the form of a multistakeholder committee, could be the curator of global principles and emerging norms for AI systems.

References

1. E. Horvitz, "AI, People, and Society," *Science*, vol 357, no. 6346, 2017, p. 7.
2. National Science and Technology Council Committee on Technology, *Preparing for the Future of Artificial Intelligence*, tech. report, Executive Office of the President, 2016.
3. S. Wachter, B. Mittelstadt, and L. Floridi, "Transparent, Explainable, and Accountable AI for Robotics," *Science Robotics*, vol. 2, no. 6, 2017; doi:10.1126/scirobotics.aan6080.
4. S. Levin, "New AI Can Guess Whether You're Gay or Straight from a Photograph," *The Guardian*, 8 Sept. 2017; www.theguardian.com/technology/2017/sep/07/new-artificial-intelligence-can-tell-whether-youre-gay-or-straight-from-a-photograph.
5. I. Rahwan, "Society-in-the-Loop: Programming the Algorithmic Social Contract," *Ethics and Information Technology*, 2017; doi:10.1007/s10676-017-9430-8.
6. P. Stone et al., "Artificial Intelligence and Life in 2030: One Hundred Year Study on Artificial Intelligence," *Report of the 2015 Study Panel*, tech report, Sept. 2016; https://ai100.stanford.edu/sites/default/files/ai_100_report_0831fnl.pdf.
7. J. McCarthy et al., "A Proposal for the Dartmouth Summer Research Project on Artificial Intelligence," 31 Aug. 1955.
8. E. Kumar, *Artificial Intelligence*, I.K. International, 2008.
9. R. Brownsword and K. Young, eds., *Regulating Technologies: Legal Futures, Regulatory Frames, and Technological Fixes*, Hart, 2008.

10. R. Calo, "Artificial Intelligence Policy: A Roadmap," *Social Science Research Network* (SSRN), 8 Aug. 2017; https://ssrn.com/abstract=3015350.
11. J. Palfrey and U. Gasser, *Interop: The Promise and Perils of Highly Interconnected Systems*, Basic Books, 2012.
12. C. Scott, "Regulation in the Age of Governance: The Rise of the Post-Regulatory State," *The Politics of Regulation: Institutions and Regulatory Reforms for the Age of Governance*, J. Jordana and D. Levi-Faur, eds., Edward Elgar, 2004, pp. 145–174.
13. R.H. Weber, *Realizing a New Global Cyberspace Framework: Normative Foundations and Guiding Principles*, Schulthess 2014.
14. U. Gasser, "AI and the Law: Setting the Stage," *Medium*, 26 June 2017; https://medium.com/berkman-klein-center/ai-and-the-law-setting-the-stage-48516fda1b11.
15. C.S. Yoo, "Protocol Layering and Internet Policy," Faculty Scholarship Paper 454, Univ. of Pennsylvania, 2013; scholarship.law.upenn.edu/faculty_scholarship/454.
16. D. Clark, "Characterizing Cyberspace: Past, Present, and Future," *MIT CSAIL*, v. 1.2, 12 Mar. 2010.
17. The IEEE Global Initiative for Ethical Considerations in Artificial Intelligence and Autonomous Systems, *Ethically Aligned Design: A Vision for Prioritizing Human Wellbeing with Artificial Intelligence and Autonomous Systems (AI/AS)*, IEEE, 2017; http://standards.ieee.org/develop/indconn/ec/ead_v1.pdf.
18. S. Abiteboul et al., "Data, Responsibly (Dagstuhl Seminar 16291)," *Dagstuhl Reports*, vol. 6, no. 7, 2016, pp 42–71.
19. National Science and Technology Council Committee on Technology, *Big Data: A Report on Algorithmic Systems, Opportunity, and Civil Rights*, tech. report, Executive Office of the President, 2016.
20. M. Scherer, "Regulating Artificial Intelligence Systems: Risks, Challenges, Competencies, and Strategies," *Harvard J. Law & Technology*, vol. 29, no. 2, 2016; http://dx.doi.org/10.2139/ssrn.2609777.
21. S. Wachter, B. Mittelstadt, and L. Floridi, "Why a Right to Explanation of Automated Decision-Making Does Not Exist in the General Data Protection Regulation," *Int'l Data Privacy Law*, 2017; https://ssrn.com/abstract=2903469.

Urs Gasser is the executive director of the Berkman Klein Center for Internet & Society at Harvard University, where he co-leads the Ethics and Governance of AI initiative, and serves as a professor of practice at Harvard Law School. His research and teaching focus on the interplay between law and technology. Gasser is a graduate of the University of St. Gallen and Harvard Law School. Contact him at ugasser@cyber.harvard.edu.

Virgilio A.F. Almeida is a faculty associate at the Berkman Klein Center for Internet and Society at Harvard University, and a professor in the Computer Science Department at the Federal University of Minas Gerais (UFMG), Brazil. His research interests include cyber policies, large-scale distributed systems, the Internet, and social computing. Almeida has a PhD in computer science from Vanderbilt University. Contact him at virgilio@dcc.ufmg.br or valmeida@cyber.harvard.edu.

This article originally appeared in IEEE Internet Computing, *vol. 21, no. 6, 2017.*

Deep Learning Triggers a New Era in Industrial Robotics

Ryo Miyajima
Preferred Networks

One reason deep learning has attracted the attention of so many researchers and engineers, even outside of the AI community, is because it can capture abstract features and recognize patterns in ways many once thought impossible for computers. The breakthrough was exemplified by the emergence of AlphaGo.[1] Prior to the successes of AlphaGo, experts had thought that the abstract strategical thoughts and theories that a human Go player develops over time through training and experience were not replicable by a computer—that is, without some technological breakthrough. AlphaGo proved that deep learning is indeed that breakthrough.

The pattern recognition capabilities of deep learning have pushed the limits in various fields—and industrial robotics is no exception. Deep learning will arguably not solve all of the problems we encounter in industrial robotics, but it will improve the perception capabilities of robotics systems, given its power to recognize complex real-world patterns robustly. Here, I examine some deep learning applications in robotics.

Automated Bin Picking

Let's first consider the application of automatically picking steel cylinders with a suction hand, as shown in Figure 1a. Given a depth camera image as input, the robot system is expected to figure out a point within the bin at which the suction hand will most reliably suck a cylinder. Conventional systems usually tackle this task by matching the image with predefined photographs or CAD data (see Figure 1b). If the object's appearance or shape cannot be predefined, conventional systems will often look for a planar surface larger than a certain area. Either way, experienced human operators must carefully tune parameters to achieve reliable results.

Researchers and engineers at Preferred Networks and FANUC have teamed up to demonstrate that deep learning could offer an alternative solution for this task. We use a depth camera image around a given suction hand position as input, together with the output representing whether that suction was successful. Through trial and error, we collect this input and output pair with the actual robot, initially starting with a random policy (see Figure 2). By training the deep neural network with thousands of these inputs and outputs, we have achieved 90 percent accuracy, which is comparable to a conventional system whose parameters are tuned by experienced operators.[2] Furthermore, our deep learning approach doesn't require us to predefine the object's appearance or shape.

Amazon Picking Challenge

Another, more complex application that might benefit from deep learning involves picking items of various shapes and forms stored on shelves. The Amazon Picking Challenge is a competition held by Amazon that aims to "strengthen the ties between the industrial and academic robotic communities to promote shared and open solutions to some of the big problems in unstructured automation" (www.robocup2016.org/en/events/amazon-picking-challenge).

For the 2016 Challenge, competitors were instructed to pick items off a shelf and place them back again. As easy as this task is for humans, the industry has yet to see automated robots replace all human pickers. Of the 15

Editor: Tao Mei, Microsoft Research Asia

(a) (b)

FIGURE 1. A bin picking application of steel cylinders: (a) Given a depth camera image as input, the robot system is expected to figure out a point within the bin at which the suction hand will most reliably suck a cylinder. (Photo courtesy of Eiichi Matsumoto; used with permission.) (b) Conventional systems usually tackle this task by matching the image with predefined photographs or CAD data. (Photo courtesy of FANUC Corporation; used with permission.)

teams that advanced to the finals, two-thirds used deep learning[3] in their system, including the top scoring teams. As an example of how deep learning is integrated into these competitive systems, I look at the work of last year's winner, Team Delft—a joint team from the TU Delft Robotics Institute and the company Delft Robotics.

I spoke with former Team Delft member Wilson Ko to learn more about how they used deep learning in their object detection component to classify objects in a camera image and output the bounding box for each object (Figure 3). Although Team Delft's success cannot be solely attributed to the use of deep learning, Ko asserted that properly integrating deep learning into conventional systems will bring us closer to fully autonomous, robust picking systems. In the following, each module in Team Delft's pipeline (the process from recognizing the object to sending a motor command to the robot) is explained in detail, based on information provided by Ko and his former teammate Mihai Morariu.

Object Detection

Because deep learning has shown impressive performance during the last few years in addressing object detection problems, Morariu said that the team decided to use it over algorithms that rely on hand-crafted features. The Faster Region-based Convolutional Neural Network (Faster-RCNN) system was one of the most popular deep neural-network-based object detection systems when the team started working on their system. It had shown state-of-the-art accuracy when it was released on datasets such as Pascal's visual object challenge (VOC) 2007 and 2012 and the Microsoft Common Objects in Context (MS COCO) dataset.[4] It also had the advantage of running at near real-time frame rates, which was essential for developing a fast robotic system. The underlying idea behind Faster-RCNN is to use a fully convolutional network that generates high-quality region proposals (that is, bounding boxes that enclose an object and their "objectness score").

Object Pose Estimation

The next step in the pipeline is object pose estimation. According to Morariu, the Team Delft system estimated the 6D pose of the detected object and matched a premade CAD model of the object against the real-time point cloud retrieved from the camera. Using the bounding box obtained from Faster R-CNN, only the relevant part of the captured point cloud was being matched. For deformable items, however, a premade CAD model is meaningless because the objects change in shape. Instead, object pose estimation was skipped and the system generated grasps directly on the filtered point cloud.

For the rigid items, Morariu explained that the team used the Super 4-Points Congruent Sets (Super 4PCS) algorithm[5] to do the matching. The system can use the transformation that is found at the end of this process to estimate the object's pose with regard to the robot. Finally, the system refines the object pose using the Iterative Closest Point (ICP) algorithm.

Grasp Generation

According to Ko, after the system determined the object's pose, it used this information to generate the grasp poses, and the grasps were predefined for the CAD model in the

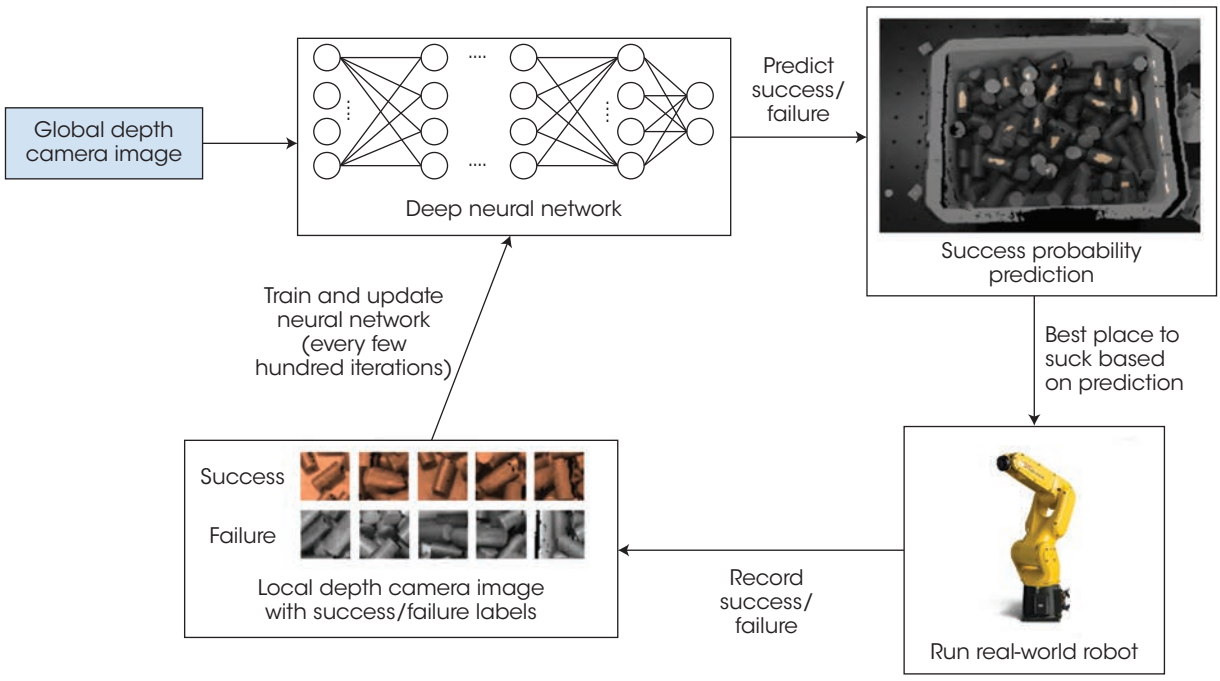

FIGURE 2. Overview of our novel bin picking system. It uses a depth camera image around a given suction hand position as input, together with the output representing whether that suction was successful. Through trial and error, we collect this input and output pair with the actual robot, initially starting with a random policy. We then train the deep neural network with thousands of these inputs and outputs.

environment. The team used shape primitives to describe the object geometry and predefine the grasps. Because an estimate of the object pose was available, the predefined grasp poses could be transformed with the object pose, providing the robot with poses to move into so it could pick up the item. The system then scored the grasps and eliminated some based on reachability and robustness. For deformable items, the system used surface normals of the segmented point cloud to generate grasp poses.

Motion Planning

Ko explained that the team distinguished between two types of trajectories: offline and online motions. Offline motions were used for motions outside of the shelf, which could be pre-generated using RRT-Connect of MoveIt (http://moveit.ros.org), based on Rapidly exploring Random Trees, RRTs, and still be collision free, because it was assumed that the environment outside the shelf would remain static.

The robot used online motions to move inside the shelf, and such motions are variable because they depend on the target's location and orientation. Online motions were split into different parts: approach, contact, lift, and retreat. Again, MoveIt was employed here for collision checking, with Trac-IK and RRT-Connect plugins for inverse kinematics and path planning.

To execute the motions, the team used the MotoROS driver from Robot Operating System (ROS)-Industrial and enhanced it to fit their needs.

Inspiring Research

Here, I introduce some promising research that will likely impact how industrial robot systems are designed in the near future.

End-to-End Training

Previously the "best practice" for designing robotics systems was to connect modular system components as a pipeline. For example, if

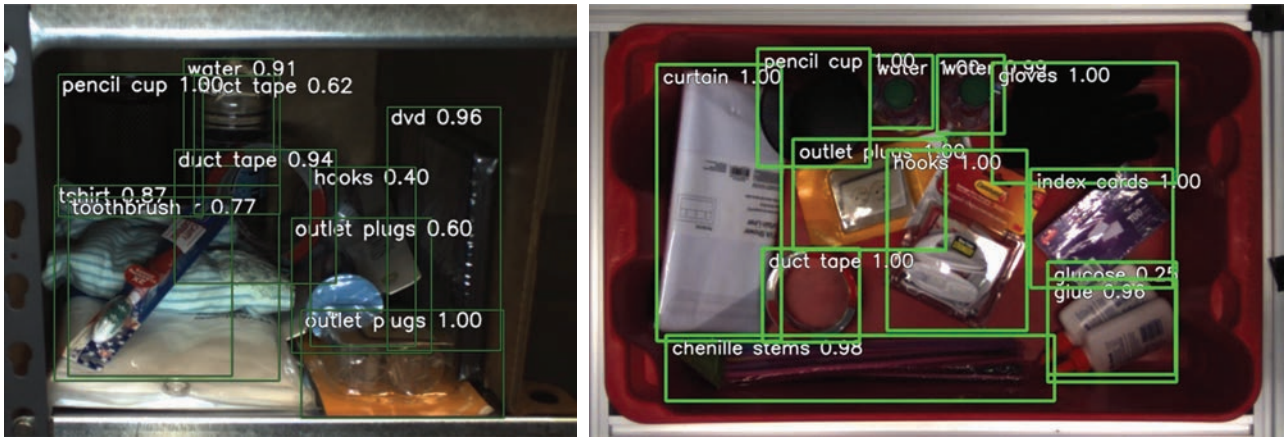

FIGURE 3. Team Delft used deep learning in their object detection component to classify the objects in the camera image and output the bounding box for each object: detection for the (a) picking task and (b) stowing task. (Photos courtesy of Carlos Hernandez Corbato; used with permission.)

you were to design a robot that serves coffee, you might build system components, such as a state estimator that realizes which phase of pouring the coffee the robot is in, a planner to designate the next action the robot should take, a controller to actuate the motors in the robot so that the desired action is achieved, and so on.

A similar approach might be applied to program a computer to play video games. Let's say you want a computer to play the game of Pong. The program receives only the video image of the game, so you could build a module to figure out the abstract state of the player (for example, where is the ball and paddle?). You could build another module to plan what the player should do (in which direction should the ball be hit back, and where should the paddle be to achieve that?), and yet another module to decide the actual commands to input into the game (which button should the player press, if any?). Finally, you'd connect all these modules.

A group from DeepMind showed that you don't need to go through the trouble of building all of these modules—by combining deep learning and some techniques from a field called *reinforcement learning*, you can directly train a computer program to output the game commands given visual images as input.[6] DeepMind's trained computer program has outperformed humans in some games.

The terminology "end-to-end" refers to methods that don't require intermediate components in this sense. Naturally enough, the transition from "pipelines" to "end-to-end" is happening in the robotics field as well. One example is the work from the University of California, Berkeley, in which a robotic arm was trained end-to-end to perform tasks like inserting toy blocks into boxes.[7]

Admittedly, with the current technology, end-to-end trained systems will not be as precise and accurate as conventional systems that are tweaked and tuned for a very specific task, such as controlling the position of a robotic arm. However, being able to teach abstract tasks (opening the cap of a bottle, for example) to robots end-to-end was not something experts thought was practical until a few years ago. With new technology, we might see robotic system designs radically reshaped in the near future.

Toward Robust Grasping

As exciting and astonishing as the state-of-the-art achievements of deep learning are, many will point out the caveats of the technology. For example, deep learning requires a large dataset for training, which is why many of those working on automated grasping have welcomed the release of Dex-Net 2.0 from the University of California, Berkeley. Dex-Net 2.0 contains

millions of datapoints to train a deep learning network[8] that tells you the quality of a grasp with a parallel jaw gripper.

Collecting large amounts of data for deep learning reminds me of the work at Google Research, where they ran as many as 14 robots simultaneously over the course of two months to collect 800,000 grasp attempts.[9] In contrast, Dex-Net does not rely on time and a vast number of robots to collect data. Rather, it exploits physics-based models so that grasping attempts can be synthesized instead of experimenting with the grasp in the real world.

For example, let's say you want to pick up a cube with two fingers. You immediately choose to place your fingers on the two facing sides and not on any other combination of sides. Your intuition about your grasp quality is in line with what grasp analytics predict. Therefore, instead of spending time and using numerous robots to predict the outcome of every possible grasp, it makes sense to use our knowledge about grasp quality and inject it into the training dataset to train a deep neural network.

One impressive trait of these methods that use deep learning is their ability to adapt to data they didn't see in the training phase—a capability is referred to as "generalization." Even without deep learning, you could program a robot arm to grasp a specific object. However, that robot will usually have difficulties adapting to other objects of various shape, friction, or appearance without reprogramming. Deep learning relaxes that restriction, and it is exactly that capability that makes it a promising method with which to tackle many of the unsolved problems and challenges in industrial robotics.

Conventional industrial robotics was all about controlling and reducing the variance of the environment so that unintelligent robots could do their repetitive work. Perhaps this practice is acceptable if the system is for mass production and the initial investment in programming/teaching the robot to perform its repetitive work will likely be recovered in the foreseeable future. However, there are two issues that need to be addressed. First, there is a growing demand for automation in mass customization solutions (as opposed to mass production). Second, even in some mass production factories that will benefit from automation, we still see human workers performing repetitive tasks that are technologically challenging (requiring dexterity, for example) or not worth the investment to automate. In either case, deep learning is a promising technology for cultivating undeveloped areas and providing us with more robust, adaptive, and reliable systems. MM

References

1. D. Silver et al., "Mastering the Game of Go with Deep Neural Networks and Tree Search," *Nature*, 28 Jan. 2016; doi:10.1038/nature16961.
2. E. Matsumoto, "Learning from 0 in a Bulk Loading Robot with Deep Learning," *Preferred Networks Research Blog*, 2015; https://research.preferred.jp/2015/12/robot_binpick_deep_learning (in Japanese).
3. "Detail Report on Amazon Picking Challenge 2016 (First Part)," *Nikkei Robotics* [in Japanese], Sept. 2016; http://ec.nikkeibp.co.jp/item/backno/RO0014.html.
4. S. Ren et al., "Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks," *Proc. Advances in Neural Information Processing Systems* (NIPS), 2015, https://arxiv.org/abs/1506.01497.
5. N. Mellado, D. Aiger, and N.J. Mitra, "Super 4PCS Fast Global Pointcloud Registration via Smart Indexing," *Computer Graphics Forum*, vol. 33, no. 5, 2014, pp. 205–215.
6. V. Mnih et al., "Human-Level Control through Deep Reinforcement Learning," *Nature*, vol. 518, Feb. 2015, pp. 529–533.
7. S. Levine et al., "End-to-End Training of Deep Visuomotor Policies," *J. Machine Learning Research*, vol. 17, no. 1, 2016, pp. 1334–1373.

8. J. Mahler et al., "Dex-Net 2.0: Deep Learning to Plan Robust Grasps with Synthetic Point Clouds and Analytic Grasp Metrics," *Proc. Robotics Science and Systems*, 2017; www.roboticsconference.org/program/papers/19.
9. S. Levine et al., "Learning Hand-Eye Coordination for Robotic Grasping with Deep Learning and Large-Scale Data Collection," *Int'l J. Robotics Research*, Mar. 2016; https://arxiv.org/abs/1603.02199.

Ryo Miyajima is an engineer at Preferred Networks. Contact him at ryo@preferred.jp.

This article originally appeared in IEEE MultiMedia, *vol. 24, no. 4, 2017.*

EDITOR: Irena Bojanova, NIST, irena.bojanova@computer.org

How Much to Trust Artificial Intelligence?

George Hurlburt, *STEMCorp*

There has been a great deal of recent buzz about the rather dated notion of artificial intelligence (AI). AI surrounds us, involving numerous applications ranging from Google search, to Uber or Lyft ride-summoning, to airline pricing, to Alexa or Siri. To some, AI is a form of salvation, ultimately improving quality of life while infusing innovation across myriad established industries. Others, however, sound dire warnings that we will all soon be totally subjugated to superior machine intelligence. AI is typically, but no longer always, software dominant, and software is prone to vulnerabilities. Given this, how do we know that the AI itself is sufficiently reliable to do its job, or—put more succinctly—how much should we trust the outcomes generated by AI?

Risks of Misplaced Trust

Consider the case of self-driving cars. Elements of AI come into play in growing numbers of self-driving car autopilot regimes. This results in vehicles that obey the rules of the road, except when they do not. Such was the case when a motor vehicle in autonomous mode broadsided a turning truck in Florida, killing its "driver." The accident was ultimately attributed to driver error, as the autonomous controls were deemed to be performing within their design envelope. The avoidance system design at the time required that the radar and visual systems agree before evasive action would be engaged. Evidence suggests, however, that the visual system encountered glare from the white truck turning against bright sunlight. This system neither perceived nor responded to the looming hazard. At impact, however, other evidence implicated the "driver," who was watching a *Harry Potter* movie. The driver, evidently overconfident of the autopilot, did not actively monitor its behavior and failed to override it, despite an estimated seven-second visible risk of collision.[1] The design assurance level was established, but the driver failed to appreciate that his autopilot still required his full, undivided attention. In this rare case, misplaced trust in an AI-based system turned deadly.

Establishing a Bar for Trust

AI advancement is indeed impressive. DARPA, sponsor of early successful autonomous vehicle competitions, completed the Cyber Grand Challenge (CGC) competition in late 2016. The CGC established that machines, acting alone, could play an established live hacker's game known as Capture the Flag. Here, a "flag" is hidden in code, and the hacker's job is to exploit vulnerabilities to reach and compromise an opponent's flag. The CGC offered a $2 million prize to the winning team that most successfully competed in the game. The final CGC round pitted seven machines against one another on a common closed network without any human intervention. The machines had to identify vulnerabilities in an opponent's system, fix them on their own system, and exploit them in

opponents' systems to capture the flag. Team Mayhem from Carnegie Mellon University was declared the winner.[2]

John Launchbury, director of DARPA's Information Innovation Office, characterizes the type of AI associated with the CGC as *handcrafted knowledge*. Emerging from early expert systems, this technology remains vital to the advancement of modern AI. In handcrafted knowledge, systems reason against elaborate, manually defined rule sets. This type of AI has strength in reasoning but is limited in forms of perception. However, it possesses no ability to learn or perform abstraction.[3]

While building confidence that future reasoning AI can indeed rapidly diagnose and repair software vulnerabilities, it is important to note that the CGC was intentionally limited in scope. The open source operating system extension was simplified for purposes of the competition,[4] and known malware instances were implanted as watered-down versions of their real-life counterparts.[5] This intentionally eased the development burden, permitted a uniform basis for competitive evaluation, and reduced the risk of releasing competitors' software into the larger networked world without requiring significant modification.

The use of "dirty tricks" to defeat an opponent in the game adds yet another, darker dimension. Although the ability to reengineer code to rapidly isolate and fix vulnerabilities is good, it is quite another thing to turn these vulnerabilities into opportunities that efficiently exploit other code. Some fear that if such a capability were to be unleashed and grow out of control, it could become a form of "supercode"—both exempt from common vulnerabilities and capable of harnessing the same vulnerabilities to assume control over others' networks, including the growing and potentially vulnerable Internet of Things (IoT). This concern prompted the Electronic Frontier Foundation to call for a "moral code" among AI developers to limit reasoning systems to perform in a trustworthy fashion.[4]

Figure 1. Some prevalent AI machine learning algorithms.

Machine Learning Ups the Trust Ante

Launchbury ascribes the term *statistical learning* to what he deems the second wave of AI. Here, perception and learning are strong, but the technology lacks any ability to perform reasoning and abstraction. While statistically impressive, machine learning periodically produces individually unreliable results, often manifesting as bizarre outliers. Machine learning can also be skewed over time by tainted training data.[3] Given that not all AI learning yields predictable outcomes, leading to the reality that AI systems could go awry in unexpected ways, effectively defining the level of trust in AI based tools becomes a high hurdle.[6]

At its core, AI is a high-order construct. In practice, numerous loosely federated practices and algorithms appear to compose most AI instances—often crossing many topical domains. Indeed, AI extends well beyond computer science to include domains such as neuroscience, linguistics, mathematics, statistics, physics, psychology, physiology, network science, ethics, and many others. Figure 1 depicts a less than fully inclusive list of algorithms that underlie second-wave AI phenomena, often collectively known as machine learning.

This myriad of potential underlying algorithms and methods available to achieve some state of machine learning raises some significant trust issues, especially for those involved in software testing as an established means to assure trust. When the AI becomes associated with mission criticality, as

is increasingly the case, the tester must establish the basis for multiple factors, such as programmatic consistency, repeatability, penetrability, applied path tracing, or identifiable systemic failure modes.

The nontrivial question of what is the most appropriate AI algorithm goes as far back as 1976.[3] The everyday AI practitioner faces perplexing issues regarding which is the right algorithm to use to suit the desired AI design. Given an intended outcome, which algorithm is the most accurate? Which is the most efficient? Which is the most straightforward to implement in the anticipated environment? Which one holds the greatest potential for the least corruption over time? Which ones are the most familiar and thus the most likely to be engaged? Is the design based on some form of centrality, distributed agents, or even swarming software agency? How is this all to be tested?

The everyday AI practitioner faces perplexing issues regarding which is the right algorithm to use to suit the desired AI design.

These questions suggest that necessary design tradeoffs exist between a wide range of alternative AI-related algorithms and techniques. The fact that such alternative approaches to AI exist at all suggests that most AI architectures are far from consistent or cohesive. Worse, a high degree of contextually-based customization is required for both reasoning and learning systems. This, of course, extends to AI testing, because each algorithm and its custom implementation brings its own unique deep testing challenges, even at the unit level.

One high-level AI test assesses the ability to correctly recognize and classify an image. In some instances, this test has surpassed human capability to make such assessments. For example, the Labeled Faces in the Wild (LFW) dataset supports facial recognition with some 13,000 images to train and calibrate facial recognition machine learning tools using either neural nets or deep learning. The new automated AI image recognition tools can statistically outperform human facial recognition capability using this dataset.[7] The task at hand, however, is fundamentally perceptual in nature. These tasks functionally discriminate through mathematically correlated geometric patterns but stop short of any form of higher-order cognitive reasoning. Moreover, while it compares selective recognition accuracy against human ability, other mission-critical aspects of the underlying code base remain unchecked under this test.

Beyond the Code

Testing machine learning becomes further complicated as extensive datasets are required to "train" the AI in a learning environment. Not only should the AI code be shown to be flawless, but the data used in the training should theoretically bear the highest pedigree. In the real world, however, datasets often tend to be unbalanced, sparse, inconsistent, and often inaccurate, if not totally corrupt. Figure 2 suggests that information often results from resolving ambiguity. Even under controlled conditions, significant differences result between the use of single or multiple well-validated datasets used to train and test classifiers. Thus, even controlled testing for classifiers can become highly complicated and must be approached carefully.[8]

Other trust-related factors extend well beyond code. Because coding is simultaneously a creative act and somewhat of a syntactic science, it is subject to some degree of interpretation. It is feasible that a coder can inject either intentional or unintentional cultural or personal bias into the resulting AI code. Consider the case of the coder who creates a highly accurate facial recognition routine but neglects to consider skin pigmentation as a deciding factor among the recognition criteria. This action could skew the results away from features otherwise reinforced by skin color. Conversely, the rates of recidivism among criminals skews some AI-based prison release decisions along racial lines. This means that some incarcerated individuals stand a better statistical chance of gaining early release than others—regardless of prevailing circumstances.[9] Semantic inconsistency can further jeopardize the neutrality of AI code, especially if natural language processing or idiomatic speech recognition are involved.

Some suggest that all IT careers are now cybersecurity careers.[10] This too has a huge implication for the field of AI development and its implementation. The question of "who knows what the machine knew and when it knew it" becomes significant from a cybersecurity standpoint. What a machine learns is often not readily observable, but rather lies deeply encoded. This not only affects newly internalized data, but—in

the IoT—these data can trip decision triggers to enliven actuators that translate the "learning" into some sort of action. Lacking concrete stimulus identity and pedigree, the overall AI-sparked IoT stimulus-response mechanism becomes equally uncertain. Nonetheless, the resulting actions in mission-critical systems require rigorous validation.

The Third Wave

Launchbury foresees the need for a yet-to-be-perfected third wave of AI, which he names *contextual adaptation*. This technology, requiring much more work, brings together strengths in perception, learning, and reasoning and supports a significantly heightened level of cross-domain abstraction.[3]

The 2017 Ontology Summit, aptly entitled "AI, Learning, Reasoning, and Ontologies," concluded in May 2017. Reinforcing Launchbury's observation, the draft summit communique concluded that, to date, most AI approaches, including machine learning tools, operate at a subsymbolic level using computational techniques that do not approximate human thought. Although great progress has been achieved in many forms of AI, the full treatment of knowledge representation at the symbolic level awaits maturity (bit.ly/2qMN0it). Correspondingly, the utility of ontology as a formal semantic organizing tool offers only limited advantages to AI and its ultimate test environment.

The semantic network involves graph representations of knowledge in the form of nodes and arcs. It provides a way to understand and visualize relationships between symbols, often represented by active words, which convey varying meanings when viewed in context. AI, largely subsymbolic today, will need to deal with applied semantics in a far more formal sense to achieve third-wave status. Under such circumstances, AI becomes nonlinear, in which cause and effect are increasingly decoupled via multiple execution threads. This leads to the establishment of *complex adaptive systems* (CAS), which tend to adhere to and be influenced by nonlinear network behavior.

Figure 2. Information provenance can often be unclear.

In a CAS, new behaviors emerge based on environmental circumstance over time. Here, there can be multiple self-organizing paths leading to success or failure, all triggered by highly diversified nodes and arcs that can come, grow, shrink, and go over time. Such networks defy traditional recursive unit testing when composed using embedded software, which is interrelated to data. This is because in a CAS, the whole often becomes far more than merely the sum of the parts.[11] Rather, new approaches, emerging from applied network science, offer a better means of assessing dynamic AI behavior that emerges over time. This becomes increasingly true as the temporal metrics associated with graph theory become better understood as a means of describing dynamic behaviors that fail to follow linear paths to achieve some desired effect.[12]

Until some reliable methodology is adopted for the assessment of assured trust within AI, the watchword must be caution. Any tendency to put blind faith in what in effect remains largely untrusted technology can lead to misleading and sometimes dangerous conclusions.

References

1. N.E. Boudette, "Tesla's Self-Driving System Cleared in Deadly Crash," *New York Times*, 19 Jan. 2017.
2. D. Coldewey, "Carnegie Mellon's Mayhem AI Takes Home $2 Million

from DARPA's Cyber Grand Challenge," *TechCrunch*, 5 Aug. 2016; tcrn.ch/2aM3iS7.
3. J. Launchbury, "A DARPA Perspective on Artificial Intelligence," DARPAtv, 15 Feb. 2017; www.youtube.com/watch?v5-O01G3tSYpU.
4. N. Cardozo, P. Eckersley, and J. Gillula, "Does DARPA's Cyber Grand Challenge Need a Safety Protocol?" Electronic Frontier Foundation, 4 Aug. 2016; bit.ly/2aPxRXc.
5. A. Nordrum, "Autonomous Security Bots Seek and Destroy Software Bugs in DARPA Cyber Grand Challenge," *IEEE Spectrum*, Aug. 2016; bit.ly/2arLOcR.
6. S. Jontz, "Cyber Network, Heal Thyself," *Signal,* 1 Apr. 2017; bit.ly/2o0ZCVe.
7. A. Jacob, "Forget the Turing Test—There Are Better Ways of Judging AI," *New Scientist*, 21 Sept. 2015; bit.ly/1MoMUnF.
8. J. Demsar, "Statistical Comparisons of Classifiers over Multiple Data Sets," *J. Machine Learning Research*, vol. 7, 2006, pp. 1–30.
9. H. Reese, "Bias in Machine Learning, and How to Stop It," *TechRepublic*, 18 Nov. 2016; tek.io/2gcqFrI.
10. C. Mims, "All IT Jobs Are Cybersecurity Jobs Now," *Wall Street J.*, 17 May 2017; on.wsj.com/2qH5VP2.
11. P. Erdi, *Complexity Explained*, Springer-Verlag, 2008.
12. N. Masuda and R. Lambiotte, *A Guide to Temporal Networks*, World Scientific Publishing, 2016.

George Hurlburt *is chief scientist at STEMCorp, a nonprofit that works to further economic development via adoption of network science and to advance autonomous technologies as useful tools for human use. He is engaged in dynamic graph-based Internet of Things architecture. Hurlburt is on the editorial board of* IT Professional *and is a member of the board of governors of the Southern Maryland Higher Education Center. Contact him at ghurlburt@change-index.com.*

This article originally appeared in IT Professional, *vol. 19, no. 4, 2017.*

Human Augmentation

Editor: Albrecht Schmidt ■ University of Stuttgart ■ albrecht@computer.org

Augmenting Human Intellect and Amplifying Perception and Cognition

Albrecht Schmidt, *University of Stuttgart*

In this first installment of *IEEE Pervasive Computing*'s new Human Augmentation department, I look at various technologies designed to augment the human intellect and amplify human perception and cognition. Linking back to early work in interactive computing, I consider how novel technologies can create a new relationship between digital technologies and humans. Forthcoming articles will provide examples of how digital systems can amplify human performance—in particular, human cognition and perception.

A REVOLUTION IN THE MAKING

It's apparent that artificial intelligence (AI) is challenging humans in many established domains. In games such as chess and GO, algorithms have outperformed humans,[1] and autonomously driven cars have started to exhibit stunning performance.[2] In both the popular press[3] and articles from esteemed colleagues,[4] we see warnings of AI posing a risk to humanity. In fact, there's a website where you can plug in your profession and job title to find whether you're at risk of being replaced by computing technologies.[5] Some experts have painted a dark future in which humans are sidelined—or even made extinct—by machines. But I disagree with these bleak outlooks!

To me, this feels like history repeating itself, and people are underestimating human abilities and flexibility. From the 18th through the 20th century, power machines—such as steam engines, combustions engines, electric motors, and hydraulic lifts—revolutionized our world, from the workplace to family life. Power (that is, muscle power) was no longer an area where humans were superior to machines. Nevertheless, these advances in physical technologies led to the world in which we now live.

In the 21st century, computing, networking, and digital media, combined with sensing and actuation, are the new ingredients for fundamental change. We're likely at the dawn of another technical revolution that will question all we know about work, economics, social environments, family, and even ethics. I suggest reading up on the first industrial revolution and its wide impact, looking in particular at the opportunities it created beyond the workplace. An interesting starting point is *Energy and the English Industrial Revolution*, by Edward Anthony Wrigley (Cambridge University Press, 2010).

THE BRIGHT SIDE OF THE FUTURE

Human history is full of technological advances that have changed how we work and live. A major talent of humans is our ability to develop tools and devices that help us adapt to different environments. The many machines invented in the last 200 years are examples of tools that have increased our physical capabilities, making humans stronger, faster, and more precise. This has inevitably changed what we value in individuals and how we structure our society. Tool use and tool making are fundamental to a species and linked to advances in evolution.

Tools for the mind and for augmenting the human intellect have been a central goal since the early days of computing (see the "Mandatory Reading: Past Visions of the Future" sidebar). Searching through vast amounts of information has become an essential tool for many professions (just try to write software without Internet access). Extending our memory and externalizing information is becoming commonplace, as media capture and access become more simplified. Tools for ubiquitous communication are providing value to people at home and in the workplace. There are many positive effects, but humans must learn how these tools fit into our lives, and it's apparent that new technologies affect how we think, sometimes literally changing our brains.[6] Evolution, however, is slow, even if the tools are quickly changing the conditions and requirements around us.

AUGMENTING THE HUMAN INTELLECT

Amplifying human abilities follows Joseph Lickider's idea of a "Man-Computer Symbiosis"[7] and extends research

 Published by the IEEE Computer Society

MANDATORY READING: PAST VISIONS OF THE FUTURE

Ideas of using information technologies to augment human cognitive and perceptual abilities have been stated by visionaries of the last century. It's exciting to go back to these powerful past visions of the future and read them with current technologies in mind.

In 1945, Vannevar Bush outlined, in "The Memex—As We May Think," a vision for making (scientific) knowledge widely available and for allowing for sharing and collaboration.[1] His vision was bold! Now, 70 years later, advances in technologies—especially in networking and capturing information—have created a world where information is readily available and where everyone can contribute with very little skill.

Joseph Licklider foresaw a close relationship between computers and humans. In "Man-Computer Symbiosis," he envisioned computers and humans working together with a high degree of flexibility, allowing for joint decision making and collaboration in solving complex (cognitive) tasks.[2] The way in which we now use computers as tools, and the seamlessness in human-computer interaction, underline that this has become the dominant way of working today.

Douglas Engelbart saw how interactive computing can augment the human intellect. He envisioned and explored experimentally how interactive applications can support humans. He coined the term of "augmenting human intellect" (see www.1962paper.org/web.html).[3] The idea was that computing technologies could increase human capability in dealing with complex problems by offering faster comprehension. Overall, he foresaw such systems solving problems that otherwise couldn't be tackled by humans.

Mark Weiser and his colleagues at XEROX PARC explored in the 1980s and 1990s how the ubiquity of digital computing, networking, and storage technologies would change our lives.[4] Technologies for networked mobile devices and large-scale interactive displays were experimentally explored from a wide range of angles, going beyond the purely technological questions, and looking at the impact on business and society.

REFERENCES

1. V. Bush, "The Memex—As We May Think," *The Atlantic*, July 1945; www.theatlantic.com/magazine/archive/1945/07/as-we-may-think/303881/5.
2. J.C.R. Licklider, "Man-Computer Symbiosis," *IRE Trans. Human Factors in Electronics*, Mar. 1960, pp. 4-11; http://groups.csail.mit.edu/medg/people/psz/Licklider.html.
3. D.C. Engelbart and W.K. English, "A Research Center for Augmenting Human Intellect," *Proc. Am. Federation of Information Processing Societies* (AFIPS), 1968; http://dl.acm.org/citation.cfm?id=1476589.1476645.
4. M. Weiser, "The Computer for the 21st Century," *Scientific American*, vol. 265, no. 3, 1991, pp. 94–104.

by Douglas Engelbart on "Augmenting [the] Human Intellect."[8] Lickider and Engelbart understood early on how interactive computing technology could be used to augment the human intellect. The idea of ubiquitous computing, as described by Mark Weiser,[9] moved the vision forward to interactive computing technologies pervasively integrated with everyday environments.

The idea of externalizing cognitive efforts and the notion of *distributed cognition* is based on the fact that user actions are situated in place and time and are often a reaction to the immediate environment.[10] One important facet of distributed cognition is that cognitive capabilities are a combination of a person's cognitive abilities and their environment. Knowledge and skill are not only in our minds but also in the environment.[11] Another important example is the intelligent use of space for problem solving.[12] By creating dynamic environments that can change technically and adapt to the tasks at hand, we inevitably alter our cognitive abilities.

Many of these technologies, ranging from early work on development tools to remote collaboration, mobile computing devices, and (more recently) machine learning, have undoubtedly augmented the human intellect for the masses. However, looking at technologies currently in the making reveals that this is only the beginning. If we can amplify ourselves, we will stay ahead of machines. Once muscle power was no longer a central need, our intellectual abilities bloomed. If machines take over basic intellectual tasks, might we have the time to achieve new abilities in social, emotional, or moral domains?

AMPLIFYING PERCEPTION AND COGNITION

The ability to see and to hear helps us make sense of our surroundings. Our perception and cognition determine how we experience the world, and human senses are incredibly powerful. However, with current and upcoming technologies, we have reached a point where sensors, databases, and algorithms can outperform human abilities compared to our raw capabilities. Imaging sensors have a higher spatial and temporal resolution than the human eye, and the amount of information we can store externally is much higher than what we can remember. Furthermore, algorithms are faster at picking out details from a vast amount of images than human observers.

The goal behind amplifying human perception and cognition is to create systems where we closely couple humans with technologies to provide us with super-human abilities. A key to such technologies is that they act like a natural extension of our own abilities, with no added effort or increased cognitive load required. Amplified perception is based on combining human senses with technical sensors. I envision new amplified senses that are seamless to use, and I'm confident they can be built with current technologies. Amplifying cognition includes a wide range of digital technologies. Examples are memory extensions and visualizations that help us take in information and enhance our short-term memory.

HUMAN AUGMENTATION

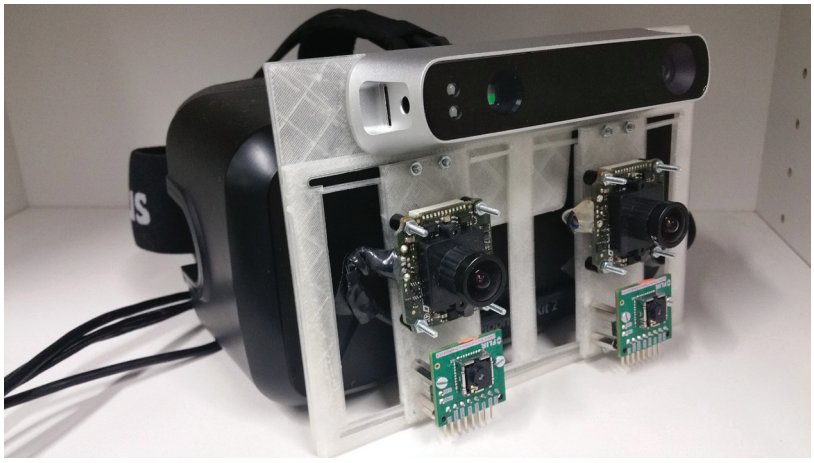

Figure 1. Prototype of a video see-through system that combines an Oculus with an RGB, depth, and thermal image and allows implicit control by different physiological sensors. (Source: Yomna Abdelrahman and Pascal Knierim, Project FeuerWeRR, University of Stuttgart; used with permission.)

Basic technologies that implement these amplifiers for cognition and perception are already available. The missing, but central, piece is the seamless integration with our existing perceptual and cognitive abilities and effective mechanisms for explicit and implicit control. Physiological sensing could be the missing link, and I expect that using eye-gaze, brain signals, and muscle activity will be key to creating a seamless user experience for amplified perception and cognition.

Amplified Perception

If we look at how to amplify perception, we see two major directions:

- enhancing and amplifying existing senses (such as vision, hearing, and touch), and
- extending perceptual abilities to domains where humans have no perception but technical sensors exist (such as sensors indicating magnetic north or solar radiation).

I expect that systems for amplifying perception could be seamlessly integrated with current perception such that, in the long term, a person wouldn't even realize the amplification.

A vision for such a system would be glasses (or contact lens, or even an implant) that let you seamlessly operate across a wide visual spectrum in which you can manipulate the focus, speed, and spectrum. Imagine you're walking along a path in the forest, and you see a squirrel in the distance. Once you look at it, you concentrate and can zoom in and see how it's nibbling on a nut. When it jumps from one tree to the next, you can, by holding your breath, slow down what you see and appreciate how it lands on the branch. Once you start walking again, you return to your normal view.

We're not there yet, but at the Human Computer Interaction Lab at the University of Stuttgart, we have experimented with different technologies for amplifying perception. Currently, technologies are still bulky—but the vision is that the technologies might be, in 20 years' time, embedded into your glasses or contact lenses. Figure 1 shows a prototype of a video see-through system that lets you have visual perception beyond the human visual spectrum. You can move between a normal color video and a thermal video, and you can add a depth view. With a similar setup, we have explored how to provide different perspectives (a first-person versus third-person view).

Amplified Cognition and Creativity

You can amplify cognition and creativity in many ways. Examples include

- amplifying a user's personal memory through contextualized capture and repeated presentation[1] (for example, with a wearable camera, a user can capture pictures throughout the day and use them for memory augmentation—see http://recall-fet.eu);
- enhancing information intake of various media (for example, for speed reading or for nonlinear viewing of videos);
- offering a parallel presentation of massive amounts of information (such as presenting large documents on large and high-resolution screens); and
- presenting related solutions to amplify creativity (such as presenting hundreds of images of existing solutions in a room).

Consider the scenario shown in Figure 2. The idea here is to build on human perceptual capabilities and support them with technologies. Instead of searching within categories, (which requires significant knowledge), we narrow down the search space to a set that humans can easily perceive and match. In this example, let's say we have 500 plants in the book. By narrowing it down to blue plants that are approximately 30 cm high, we could probably reduce the set to 50 plants. We assume that a person could quickly pick the right plant from a set of just 50, shown on a high-resolution wall-sized display. By not fully automating the task, and by adding a human perceptual step, we can increase the person's knowledge as she sees what's around the actual match.

Quantifying Cognitive and Perceptual Amplification

Even though there are now many tools that augment and amplify our abilities, solid metrics don't exist, and we, as a community, have made little effort

to quantify such augmentations and amplifications beyond single cases. For tools in the physical world, the effect is typically easy to quantify. A human without any technical support can travel about 5 km an hour; with a bike, that goes up to 20 km per hour, and with a car, up to 100 km. You can also clearly quantify the difference in using an electric drill versus a mechanical drill, noting the number of holes drilled and the increase in precision.

In the digital world, there's little scientific work on quantification. How much quicker can you develop software with a 50 mbit/s Internet connection versus a 1 mbit/s connection? What is the effect on the productivity of a software developer if she has access to the website stackoverflow.com compared to using a printed reference manual for the programming language? Is having an interactive 60-inch screen in the meeting room more effective than everyone having a tablet computer? How much quicker are you at solving problems when you can access YouTube versus using the manual provided with the product?

To scientifically validate the amplification of human capabilities, it's essential to understand how to quantify the amplification. Can we create metrics that let us state that using a certain digital tool will increase your externally perceived IQ by 5 points? Or that using another tool will allow you to perform (with the tool) as well as someone (without a tool) whose average school grades are 0.5 points better than yours?

Validating amplifications obviously won't be easy, and the methods will be disputed. Furthermore, validation will most certainly require complex, large-scale experiments, but such experiments are essential to making advances not just visible but also measurable.

Figure 2. A sample scenario—naming a plant. This examples illustrates comparing the usage of a traditional book based on categories with a large screen presentation. All potential candidates are shown at once, but filters (based on color, size, and so on) can be applied. Spotting the match relies on human perceptual abilities. (Artwork by Katrin Wolf; used with permission.)

TOWARD SUPER HUMAN TECHNOLOGIES?

It's apparent that many of these ideas will also question what we really want humans to be. Do we create (or at least attempt to create) technologies that make us smarter than the technology, or do we assume that technologies will take over?

Many are skeptical of technologies that amplify our abilities. There's a natural skepticism about creating super humans through technologies, but looking back at the machines that replaced human muscle power shows us that super human technologies aren't new. Humans move at great speeds (in cars) and lift amazing weights (with cranes). Technologies have and always will change our abilities. The introduction of written texts, book printing, and photography are just some examples that made us super human with regard to our cognitive abilities—compared to people who don't have these technologies. We also have perceptual aids (such as microscopes or thermal cameras) as well as cognitive aids (such as calculators).

So what's new? The difference now is that the upcoming technologies for cognition and perception are moving much closer to our bodies. Natural and implicit control will thus make us feel as though such technologies are a part of us—if we get it right.

ACKNOWLEDGEMENT

This work has received funding from the European Research Council (ERC), under the EU's Horizon 2020 research and innovation program (grant agreement no. 683008) and from the German Federal Ministry of Education and Research (BMBF) in project FeuerWeRR.

REFERENCES

1. C. Moyer, "How Google's AlphaGo Beat a Go World Champion," *The Atlantic*, 28 Mar. 2016; www.theatlantic.com/technology/archive/2016/03/the-invisible-opponent/475611.

HUMAN AUGMENTATION

2. J. Vincent, "World's First Self-Driving Taxi Trial Begins in Singapore," *The Verge*, 25 Aug 2016; www.theverge.com/2016/8/25/12637822/self-driving-taxi-first-public-trial-singapore-nutonomy.

3. M. Sainato, "Stephen Hawking, Elon Musk, and Bill Gates Warn about Artificial Intelligence," *The Observer*, 19 Aug. 2015; http://observer.com/2015/08/stephen-hawking-elon-musk-and-bill-gates-warn-about-artificial-intelligence.

4. S. Hawking et al., "Transcendence Looks at the Implications of Artificial Intelligence—But Are We Taking AI Seriously Enough?" *The Independent*, 1 May 2014; www.independent.co.uk/news/science/stephen-hawking-transcendence-looks-at-the-implications-of-artificial-intelligence-but-are-we-taking-9313474.html.

5. Q. Bui, "Will Your Job Be Done by a Machine?" *NPR*, 21 May 2015; www.npr.org/sections/money/2015/05/21/408234543/will-your-job-be-done-by-a-machine.

6. N. Carr, *The Shallows: How the Internet Is Changing the Way We Think, Read and Remember*, Atlantic Books, 2010.

7. J.C.R. Licklider, "Man-Computer Symbiosis," *IRE Trans. Human Factors in Electronics*, Mar. 1960, pp. 4–11; http://groups.csail.mit.edu/medg/people/psz/Licklider.html.

8. D.C. Engelbart and W.K. English, "A Research Center for Augmenting Human Intellect," *Proc. Am. Federation of Information Processing Societies* (AFIPS), 1968; http://dl.acm.org/citation.cfm?id=1476589.1476645.

9. M. Weiser, "The Computer for the 21st Century," *Scientific American*, vol. 265, no. 3, 1991, pp. 94–104.

10. L.A. Suchman, *Plans and Situated Actions: The Problem of Human-Machine Communication*, Cambridge Univ. Press, 1987.

11. J. Hollan, E. Hutchins, and D. Kirsh, "Distributed Cognition: Toward a New Foundation for Human-Computer Interaction Research," *ACM Trans. Computer-Human Interaction* (TOCHI), vol. 7, no. 2, 2000, pp. 174–196.

12. D. Kirsh, "The Intelligent Use of Space," *Artificial Intelligence*, vol. 73, no. 1, 1995, pp. 31–68.

Albrecht Schmidt is a professor at the University Stuttgart, Germany. Contact him at albrecht.schmidt@vis.uni-stuttgart.de

This article originally appeared in IEEE Pervasive Computing, *vol. 16, no. 1, 2017.*

Editor: **Robert Blumen**
Salesforce
robert@robertblumen.com

Katie Malone on Machine Learning

Edaena Salinas

MACHINE LEARNING WAS featured in episode 193 of Software Engineering Radio with Grant Ingersoll in 2013. But because this area has changed considerably in the past four years, it made sense to revisit it with a fresh outlook. In episode 286, Edaena Salinas talks with Katie Malone, a data scientist in the R&D department at Civis Analytics, which specializes in data science software and consulting. Katie earned a PhD in physics from Stanford University; during her studies she searched for new particles at CERN. She teaches Udacity's Intro to Machine Learning course and hosts Linear Digressions, a podcast about machine learning (lineardigressions.com).

Here, Katie and Edaena discuss the major types of machine-learning algorithms and some examples, including supervised and unsupervised classification. Portions of the interview not included here for reasons of space include topics such as cleaning the raw data, training data versus test data, randomization, evaluation metrics, and Katie's take on popular programming languages. To hear the full interview, visit se-radio.net or access our archives via RSS at feeds.feedburner.com/se-radio. —*Robert Blumen*

Edaena Salinas: Machine learning is widely used—in search engines, speech recognition, language translation, Netflix recommendations, and most recently in driverless cars. In the coming years, we'll see it used in more fields. So, what is machine learning?

Katie Malone: My background is in science. I believe that there's truth in the world and that science is one of the ways we get to that truth. It's really hard to measure truth directly. Instead, we collect data on the world. If we analyze that data, sometimes we can pull out the truth. A true thing about the world might be, "I'm interested in watching this movie." Or it might be, "There's a good way to translate this sentence from English to French." Machine learning is, in my view, a suite of tools that allows you to analyze data to figure out what's going on in the world, and how that's expressed in the data.

Usually it involves heavy computational lifting. The "machine" component implies computers. Then there's usually a heavy dose of statistics, and often additional scientific fields. If you're studying human behavior, you should be aware of [other fields that study humans] like behavioral psychology and economics. Those areas give you context about the thing you're interested in.

Edaena: How does machine learning relate to AI?

 Published by the IEEE Computer Society

Katie: I once heard that machine learning focuses more on understanding—measuring or making predictions—while AI is thinking one step further. Once we understand what's going on, how can we make better decisions? How can we change the way we do things to take advantage of those insights? AI adds a layer of decision making on top of machine learning.

Edaena: Let's walk through a simple example of machine learning: spam detection in email. Once I indicate that an email is spam, I'm telling the system something. What happens under the hood?

Katie: Email is an example of supervised classification. Let me break this into two parts. Supervision occurs when you have the correct answer for some of the cases. In this example, you provide the answer when you manually label the email as spam. If you don't tell the model that it's spam, then the model assumes it's a legitimate email. Classification is sorting things into two buckets: spam and not spam.

Machine learning is making predictions based on the attributes of an email. We learn what spam email looks like, and then we extrapolate those patterns onto new emails to predict whether they're spam or not.

The model is probably going to look at the words in the email, and potentially the sender's domain. Spam emails have very particular patterns. The words in spam tend to be distinctive—usually they're trying to sell you something with lots of superlative adjectives. Or maybe they're trying to get you to send money to somebody in a foreign country. Very often there are grammatical mistakes.

Based on the presence of particular words like "Nigerian prince," from the cases where you have said "this is spam," the model can learn those patterns and apply them to new cases. Hopefully at some point you don't have to manually label emails because the model will have figured out what spam looks like.

Edaena: By that time, is the system able to figure out those common words?

Katie: Spam is an interesting case, as presumably spammers are getting more sophisticated. The spam filters that worked five years ago probably wouldn't work that well right now. That's another important aspect of machine learning: it's pretty rare to have a problem that you solve once and for all. Usually you want to revisit it periodically to see if the solutions you came up with last year or last month still apply.

Spam is a good example of that. I don't know if people talk about Nigerian princes anymore, because that's such a cliché at this point, but the formula of "We're going to pretend there's money sitting in an account and if you send a small deposit we'll release it to you" remains popular, although the exact details change. In that scenario, you have to keep retraining your algorithm to continue to make good decisions.

Edaena: How is this information represented? Is there a specific format for the model?

Katie: The simplest thing you can do is to treat each word as its own feature. Many machine-learning algorithms assume there's a big matrix of attributes. Imagine a matrix as a big data table, and each row of the table is an email and each column is a word. If a particular word shows up in a particular email, you'll get a 1 in that spot in the matrix; if it doesn't show up, you'll get a 0. Then you can put that matrix into a standard machine-learning algorithm, and it will find the structure in the matrix that allows it to understand which words are most closely associated with the emails you've classified as spam.

Another important aspect of machine learning is thinking about different representations of your data. What I just described is the simplest way you might represent data in an email, but there are other algorithms that can be more compressed with respect to how the words or sentences showing up in an email are represented. How you represent your data has an intimate connection with the type of algorithm you're going to use to do the supervised classification. The way the data is formatted can make it very easy for us to find the truth we're seeking, or it can make it very hard. It's worth thinking about carefully.

Edaena: What's one way in which this data has been compressed in other data structures?

Katie: One case is Netflix movie recommendations. Imagine that each person who watches movies is a row and each possible movie they could watch is a column. Most people will only watch one percent of all the movies out there. And most movies are not going to be watched by even a significant fraction of all users. You have a big sparse matrix.

In a case like this, you can use matrix factorization. Instead of having this big sparse matrix, imagine that there are two factors—two different types of attributes—that we're trying to understand. [In this case, users and

SOFTWARE ENGINEERING RADIO

Visit www.se-radio.net to listen to these and other insightful hour-long podcasts.

RECENT EPISODES

- 283—Host Felienne talks with Alexander Tarlinder about developer testing and his book on the topic.
- 289—James Turnbull returns to the show to tell host Robert Blumen how to automate infrastructure builds using declarative programming with the Terraform tool.
- 290—Docker's Diogo Mónica talks with host Kim Carter about Docker security aspects.

UPCOMING EPISODES

- 292—Philipp Krenn talks to host Jeff Meyerson about the search server Elasticsearch.
- 293—Yakov Fain and new host Matthew Farwell discuss the popular JavaScript framework Angular.
- 294—New host Alex Newman talks asynchronous I/O with Rust language expert Carl Lerche.

types of movies.] Say there are buckets, or segments, of users, and users in each bucket watch certain types of movies or certain mixtures of movies. Then we have types of movies: action movies or foreign documentaries. Whether a particular user likes a particular movie is a combination of the type of user and type of movie. Representing the same data differently can make it easier and more direct to figure out whether a user is going to like a movie.

Edaena: How does machine learning handle cases where there's no prior data?

Katie: In recommendation engines, there's the classic problem of "cold starts." This is when a new movie will be added to Netflix next month, and they need to figure out if a lot of people are going to want to watch it. Should they give it valuable real estate on the front page to advertise it? And to which people? But they have no data on this movie yet. They don't know who has watched it or who liked it before. This is a tricky place to be in terms of machine learning, because machine learning is usually about pattern recognition, and there's no pattern yet.

But if you have some other contextual information, like "this is an action movie," then you have a better place to start from. You have some idea of the people who like action movies. This is effective at the beginning, and then you can refine those estimates as you collect more data.

Edaena: Can supervised learning be applied to values that are continuous instead of discrete?

Katie: That's usually called regression. A lot of the same algorithms can be used for classification and regression, depending on the final type of output you want.

Edaena: Can you explain regression a bit more? For example, what is the objective of linear regression?

Katie: Linear regression has a continuous output. Classification is trying to figure out if something is A or B, spam or not spam. You wouldn't say that spam has an inherently higher value than not spam or vice versa. There isn't a natural ordering of those two things.

We'll use income as an example. If you're trying to predict somebody's income from other attributes that you have, then obviously there's a natural ordering. There's a natural ordering to values like $10,000 and $100,000.

Linear regression tries to use known attributes about a person to predict another attribute, like income. Do I see a relationship between a person's attributes and their income? One example is age. The older someone is, up to a certain point, the more money they tend to make. You might observe what kind of car this person drives. I know there are patterns in income versus the type of transportation you use: richer people have nicer cars.

None of these patterns is going to hold absolutely for every single case, but statistically it'll usually hold. And from that you can make predictions. The quality of those predictions will depend on how good your data is, and to a lesser extent, how good your algorithm is. But you're getting a little closer than if you were to make a shot-in-the-dark guess.

Edaena: If we plot the income on one

axis and the car's price on another, you could predict from income how much you spend on a car, right?

Katie: That would give you some idea. You could imagine fitting a line to the distribution you see. The slope of that line will give an estimate of someone's income once you know how expensive their car is. For every data point in your dataset, the line gives a predicted income.

Edaena: How do we measure how good the line is compared to the data?

Katie: For a lot of machine-learning algorithms, there are standard metrics. Somebody drives an expensive car, so I think they make $100,000 a year. If you have the actual income in your dataset, then the dataset shows that the person makes $110,000 a year. The $10,000 difference is my error. Sometimes you square it for other reasons that we don't have to get into. You sum that over your whole dataset and divide by the number of points. This is a metric for the "goodness of fit" of this line. For classification you can use accuracy, which is how many you got right divided by the total.

One of the big challenges of machine learning is figuring out if those metrics are really measuring the thing you care about, because usually they aren't. You have to be a little smarter to figure out exactly what you care about, and modify the metrics to reflect that. You usually have good options, but the artistry is knowing when to leave those options behind.

Edaena: How is unsupervised machine learning different from supervised?

Katie: The canonical answer is that in supervised machine learning, you have correct answers that came with your dataset. With unsupervised machine learning, you don't have that luxury. You just have data. The types of questions you can ask of that data are different and very often constrained by the fact that there's no correct answer.

There are different types of questions you try to answer with unsupervised learning. In my experience, it's really hard. It's a different way of thinking about your data, because when you have the correct answer, you want to get as close as possible to it. With unsupervised machine learning, it can be very tricky to try to understand what a good answer looks like.

Edaena: What do some of those questions look like?

Katie: The biggest one in my experience has been clustering. Clustering is the idea that you have blobs or coherent clumps in the data. You want to find them. This is hard because with a lot of real-world datasets, you don't know if there are clusters in the dataset to begin with. If you don't find any, it's really hard to know if it's because you're doing a bad job or because it doesn't exist. This gets back to the idea of truth—what are we trying to understand here?

Principal component analysis is another unsupervised technique. You're trying to find ways

This article originally appeared in IEEE Software, *vol. 34, no. 4, 2017.*

of compressing the data down to the aspects that make it the most variable. You're trying to find directions in your dataset in a lower-dimensional space that maximize the variance in your data.

The point is that this is something you can do without knowing any "correct" answers for your data, which can be pretty useful when you don't have labels to rely on.

Edaena: Is unsupervised machine learning widely used? Are there any systems that we interact with that might be using it?

Katie: Segmentation within marketing is one example. If you think your users might fall into a few main buckets, you want to segment them. For example, buyers of computers might be power users, programmers and data scientists, people who watch movies and use Facebook, and people who use it as a work machine but don't program on it. Maybe there's a dataset that would allow you to pick out these distinct groups.

Edaena: Unsupervised might also be more about discoverability, because you don't necessarily know what you're looking for.

Katie: That's fair. Supervised methods are used when you know what you're trying to answer. Unsupervised methods are for when you don't have that labeled data available or when you don't exactly know how to slice and dice the data yet.

EDAENA SALINAS is a software engineer in Microsoft Research's Knowledge Technologies Group. She also hosts The Women in Tech Show (thewomenintechshow.com). Contact her at edaena@thewomenintechshow.com.

EDITOR: Seth Earley, Earley Information Science, seth@earley.com

The Problem With AI

Seth Earley, *Earley Information Science*

Science is actually pretty messy. When I was a chemistry undergraduate, I loved the theory behind biochemistry—the endless complexity allowed by simple rules; how massive, complex cellular machines could arise from a few building blocks. In lab, however, I struggled to make the simplest reactions work. Starting with pure crystalline compounds and expensive laboratory equipment, when the result was also expected to be crystalline, I ended up with piles of brown goo—with my instructor concluding, "Well, it could be in there" in reference to the experiment's objective.

Data science is also very messy. Frequently the starting point is the data equivalent of brown goo—messy, poor quality, inconsistent data—with the expectation that pure crystalline results will be the output of the next best action, personalized marketing campaigns, highly effective custom email campaigns, or a cross-department, cross-functional, 360-degree understanding of customers and their needs.

Artificial intelligence (AI), though broadly applied these days to mean almost any algorithm, is primarily dependent on some form of machine learning. Machine learning in turn is frequently fueled by what is called big data (high-velocity, high-volume, highly variable data sources) but can also be fueled by traditional data sources.

Variable Does Not Mean Poor Quality

There is a common misconception that "variable" data can mean "messy" data and that "messy" data can mean "poor-quality" data. Simply put, variable does not mean messy, and messy does not mean poor quality. Variable data is data that has different formats and structures. To use it, we need to understand how the different types of data can be used as signals to achieve a result. Twitter data is very different than transactional data. The two together can provide insights about how social trends impact sales. Messy data can be missing values or can be in formats that are difficult to ingest and process. The data can be very good, but requires work to get it into a format for processing.

A recent article in *Sloan Management Review* stated that

> *Organizations can now load all of the data and let the data itself point the direction and tell the story. Unnecessary or redundant data can be culled … [This process is] often referred to … as 'load and go.'*[1]

While conceptually accurate, there is much left open to misinterpretation. "All the data" needs to be defined. Does it mean all product data, social media data, accounting data, transactional data, knowledge base data? Clearly "all" is an overgeneralization. And this approach has its drawbacks. Sandy Pentland, MIT professor, remarked at the recent MIT CIO Symposium that "Putting all of your data in a data lake makes it convenient for hackers to go to one place to steal all of your data."

No matter what the scope is, we have to select data that is appropriate to the domain of the problem space. The data needs to be in a consistent format. It cannot contain incorrect values. If the data is incorrect or missing, then the algorithm cannot function correctly unless we are making accommodations for those issues. "It's an absolute myth that you can send an algorithm over raw data and have insights pop up," according to Jeffrey Heer, a professor of computer science at the University of Washington, as quoted in the *New York Times*.[2]

Technology writer Rick Delgado notes that "many data scientists jokingly refer to themselves as data janitors, with a lot of time spent getting rid of the bad data so that they can finally get around to utilizing the good data. After all, bad data can alter results, leading to incorrect and inaccurate insights."[3]

In a recent conversation I had with Laks Srinivasan, chief operating officer (COO) of Opera Solutions, he asserted that "80 percent of the work the data scientists are doing is data cleaning, linking, and organizing, which is an information architecture (IA) task, not a data scientist function."

Opera, founded in 2004, was one of the firms that tied for first prize in a Netflix contest that was offering US$1 million to the company that could beat its recommendation engine by 10 percent or more. (The three-year contest, which ended in August 2009, awarded the prize to a team from AT&T Labs, which submitted its response just minutes before Opera.) Opera is an example of a company that developed a platform to help data scientists in many aspects of analysis, feature engineering, modeling, data preparation, and algorithm operationalization.

A Range of AI Applications

AI applications exist along a spectrum. At one end lies embedded AI, which is transparent to the user, but makes applications work better and easier for them. Spelling correction is an example that people take for granted. Machine translation is another. Search engines use machine learning, and AI, and, of course, speech recognition, which has made enormous progress in recent years.

At the other end of the spectrum are the applications that require deep data science and algorithm development expertise. The people who develop these applications are technical experts with deep mathematical and data science knowledge. They devise and tune the algorithms that provide advanced functionality.

Along the continuum are the platforms and development environments that make use of the tools (many of which are open source). These applications require various levels of configuration and integration to provide capabilities.

Types of Cognitive Computing

For example, consider a type of "cognitive computing" application. Cognitive computing is a class of application that helps humans interface with computers in a more streamlined, natural way. Such applications are also capable of processing information in a less traditionally structured manner to provide a range of answers, with probabilities based on the user's context and details about the data sources.

One type of cognitive computing application is the processing of large amounts of patient observational data and providing a "second opinion" about a diagnosis. Physicians are using this approach to augment their knowledge and experience when developing treatment regimens. Another type is creation of an intelligent virtual assistant (IVA) that retrieves answers to procedural questions rather than lists of documents. IVA functionality requires various mechanisms that are powered by machine learning. The first is speech recognition, which translates spoken language into text. The next is a mechanism for deriving intent from the user query or utterance. Intent can be based on training sets and examples of phrase variations, or it can be from parsing language to derive meaning.

The Role of Machine Learning

Each of these approaches leverages machine learning. Some dialog management approaches can use mechanisms akin to language translation. Given enough questions and enough answers, a machine learning algorithm can "translate" questions into the correct responses. When the intent is derived via natural language understanding or training set classification, a response can be retrieved from a corpus of content via a ranking algorithm that uses signals generated through determining the intent of the user as well as additional metadata that can inform the user's context—anything from purchased products, to configured applications, to demographic or social media data.

Inference can use relationships mapped in an ontology—for example, products associated with a particular solution or steps to troubleshoot a specific device configuration. Some of this knowledge is inferred from the data and some is intentionally structured—the knowledge engineering approach to AI.

Contextualizing Endless Knowledge Sources

Organizations have enormous repositories of knowledge in the form of processes, procedures, manufacturing techniques, research methodologies, embedded designs, programming code, configured applications, technical documentation, knowledge bases of various kinds, engineering libraries, expert systems, traditional libraries, technical publications, scientific, engineering, and trade journals—the list of explicit knowledge sources is endless. Historically, humans have always limited the scope of the information that they consume—for example, by picking up a book on a topic, searching for a specific area in a library, pursuing a specialized library, or seeking out a particular journal. Even in our digital age, engineers will go to engineering sites for nuanced, specialized information. Scientists will go to scientific sites, and so on.

Information from highly diverse sources cannot be processed as raw data inputs for any purpose without restriction. It needs to be parsed, curated, packaged, contextualized, and componentized for consumption by users or ingested by systems for application to a limited number of scenarios. As powerful as it was, the Jeopardy-playing Watson program required specific information sources to function correctly.

Can Curation Be Automated?

Machines can help when given the correct scaffolding and representative training sets. Data and content sources can be processed by machine algorithms, overlaying the structure and identifying patterns in the information to assist in componentization and contextualization. The process is iterative and requires human judgment and inputs to fine-tune results. Those results might be the componentized information containing specific answers to questions rather than large amounts of text. When the content is fine-tuned and componentized, the specific answers can be more readily retrieved. A user looking for an answer does not want a list of documents, but the answer to the question. Bots and intelligent virtual assistants are designed to respond with an answer or a short list of suggestions presented in the correct context (the user's query or intent). Autotagging and autoclassification machine learning algorithms can apply the correct metadata to content to allow for those contextualized results.

The Role of Ontologies

Ontologies are the containers of metadata—the knowledge scaffolds or structures that can be abstracted from systems of knowledge and applied to other bodies of information for organization and contextualization. The ontology can capture the relationships between knowledge elements and ways of organizing those elements—for example, the list of user intents with corresponding actions. A taxonomy of products can be related to a taxonomy of solutions composed of those products. Or a list of problem types can be associated with corresponding troubleshooting approaches.

Tools such as virtual assistants become channels for knowledge structured with an ontology, along with rules and contexts that apply to specific problem sets. Take, for example, the task of servicing a customer who is trying to set up and operate a new fitness tracker. Instead of searching on the website or calling the help desk, the customer might try typing a question into the company's support chat bot. The bot interprets the natural language question as an intent, and the ontology allows retrieval of the correct responses from a knowledge repository. The ontology manages intents and responses as well as terminology and phrase variations for algorithm training.

The advantage of a natural language question over a search is that it becomes easier to derive the user's intent when they ask a fully formed question rather than typing a few ambiguous keywords. A bot can also be programmed to further disambiguate intent by requesting more detail from the user. This type of natural interface can also be used to access corporate information sources—running a financial analysis or retrieving underwriting procedures, for example.

Maturing Algorithms Still Necessitate Data Clean-Up

While machine algorithms play an important role in both the preparation of data and interpretation of user intent, these types of applications require a significant amount of knowledge engineering to be successful.

As machine learning algorithms mature, the heavy lifting will become more invisible and behind the scenes, and data or content preparation as well as application tuning and configuration will constitute the bulk of the work and require the greatest effort. With data scientists increasingly in short supply, business users will need to perform more analysis so that a backlog does not develop behind scarce data science resources. Data preparation is a major challenge, and operationalizing capabilities is an even bigger one. This is because knowledge of deep analysis approaches is becoming lost in translation from the laboratory

environment to the operational environment. Given that detailed machine learning approaches are less accessible to business people, there is increasingly a gulf between the business world and the IT world. However, two trends are in play. Sophisticated tools are becoming more commoditized, while more advanced capabilities are being made available to business people through platform approaches. The key component of data preparation, data operationalization, and translation between business challenges and analytical tools is the semantic layer—the glossaries, thesaurus structures, metadata standards, data architectures, and quality mechanisms.

As the tools get more mature, organizations will get value from them only if they take control of the things that will not be commoditized by the marketplace—their data, content, processes, and semantic translation layers. For example, organizations will not get a competitive advantage by building speech recognition. That problem has been solved (for the most part—it is still improving, but building the algorithms from scratch would not have business value). They will, however, gain a competitive advantage from servicing their customers uniquely with a speech recognition agent that accesses the knowledge they have about their customers and serves up the products and content they need.

Rethinking High-Power Analytics

As demand is exploding for big data analytics, data scientists are increasingly in short supply. When a company is building predictive models or machine learning models, a few factors stand out.

Every journey starts out with raw data, so if a company is doing multiple projects for the same client and the same department, multiple teams start with the same raw data, which can be inefficient. The second factor is that so much of the work data scientists are doing is data cleaning, linking, and organizing, which, as Srinivasan mentioned, is an IA task, not a data scientist function.

The third factor is that even after the data is cleaned up and models are developed that accurately predict (for example) who is likely to buy a certain product, it takes a lot of time to go from the data science sandbox to actually operationalizing the analytics that create a business impact.

This disconnect occurs because the development environment and the production environment are very different. As Srinivasan explains, "The data scientists might build a model using SAS in the sandbox and using certain datasets, but the IT department needs to re-code the variables and models in Java or optimize R code to scale in Hadoop when the application goes into production. At this point, the data is also very different because it goes beyond the test datasets, so the data scientists have to retest it against the model. Finally, even when the projects are in production, all these insights and know-how [are] fragmented into documents or code or people's heads. As staff turns over, knowledge is lost."

Re-Imagining the Analytics Lifecycle

When Opera began considering how to address these issues, it came up with the approach of fundamentally re-imagining the analytic development lifecycle by developing a "semantic layer" between the data layer (raw data) and the use case, application, and UI layer. The thought was that the company could preprocess the data to a point, independent of its future use, and then apply AI and machine learning in converting big data to small data. By putting a semantic layer around analytic models and tools, all the users can find them once the semantic layer is operationalized.

According to Srinivasan, "By making data independent of use cases and operationalizing it, and then making it machine-learning-driven and AI-driven, the signals learn about the data. The system becomes a learning system, not a static, one-time data modeling system. It becomes a continuous feedback, loop-based, living, breathing kind of a central nervous system, in the enterprise."

In other words, the semantic layer acts as a way to translate business problems into the inputs needed to query a big dataset. The technical predictive algorithms operate under the covers, and this complexity is hidden from the user. The algorithms simply have to point to the big data sources (that are correctly cleansed and prepared, of course) and then provide their parameters as inputs to predict their outcomes, run simulations, segment audiences, customize campaigns, and so on.

Developing an Orchestration Layer

In the case of Opera, the company went on to build a platform from the ground up to create and manage the signal layer, and ran mission-critical applications on it. The platform, called Signal Hub, processes data from about 500 million consumers for global blue-chip clients across industries. This approach allowed Opera to essentially outsource the data science work, operate on its platform, and sell solutions to business buyers. When Opera developed and then productized

This article originally appeared in IT Professional, *vol. 19, no. 4, 2017.*

the platform as an orchestration layer in 2013, many organizations did not have the IT or data science resources to fully exploit the power of advanced tools. The market has matured since then, and that strategy—to productize as an end-to-end AI and machine learning enterprise platform by hardening with security, scalability, and governance capabilities—provides valuable lessons for organizations building data-driven solutions.

Thinking about data as a service and the platform as an orchestration layer between business problems and technology solutions can help organizations achieve dramatic improvement in data scientists' productivity, and in the productivity of business analysts and business intelligence workers. "The maturing of technologies and emergence of platforms is democratizing insights derived through machine learning and capabilities provided by AI in a way that we say makes ordinary people extraordinary," says Srinivasan. "If all the insights and expertise are buried in a small team within a company, it doesn't really leverage the value of AI tools to be used by an average call center rep."

The concepts of data as a service and platforms as an orchestration layer have far-reaching implications for the future of AI-driven enterprises. Not only can data be more fully exploited by this paradigm, but so can knowledge and content—the raw material on which cognitive applications are being developed. According to Henry Truong, CTO of TeleTech, a $1.4 billion call center services firm, "Organizations can normalize knowledge in the same way that they normalize data—through componentizing knowledge into the building blocks that provide solutions to problems. The knowledge ontology becomes the data source to orchestrate more and more process actions, that, in our case, prevents service disruptions." This approach is beginning to be exploited in ways that allow for interoperability between platforms that are exposing functionality through a services layer. Those "normalized knowledge bases" are powering chat bots that are driving the next-generation digital worker.

Leveraging Platforms and Orchestration Layers

Many organizations are attempting to build their own platforms and believe this is required to create a competitive advantage from machine learning and AI capabilities. The key decision point is whether the platform is the differentiator or whether it is the data and orchestration layer that will be the differentiator. "I frequently hear CIOs say they have a platform or that they are building machine learning. The problem is that it is easy to go through $100 million or more, and a lot of pain and suffering. I say, 'Do not try this at home' in my presentations and hope they take it to heart," cautions Srinivasan.

A core premise for success with advanced analytics is that organizations need to build metadata structures and ontologies to define relationships among data elements relevant to their companies. Srinivasan continues: "That is the investment that organizations should be making rather than building their own platforms. They should be building their own representation of the core of the business, the soul of the business, which is the ontology that can embody all that knowledge of processes and customers. Insights can then be fed back into the ontology, so it becomes that living, breathing thing. It is a semantic layer that evolves around that."

Most of the work that data scientists do is "data janitorial" work, as opposed to science work, and there is a gulf between prototype and sandbox, and innovation and production. In addition, having pockets of knowledge and expertise throughout the enterprise, which may be gone when an employee leaves, poses a problem when the knowledge is not institutionalized or captured in a system. Organizations are best off if they focus on understanding their own data, focus on the business problems they are trying to solve, and build the semantic layers that can allow for data portability across various platforms. This lets them take advantage of best-of-breed solutions and not become locked into a particular vendor that does not abstract the business problem, analytic, data, and platform layers required to operationalize the fast-evolving advanced machine learning analytic and AI technologies. IT

References

1. R. Bean, "How Big Data Is Empowering AI and Machine Learning at Scale," *MIT Sloan Management Rev.*, 8 May 2017; bit.ly/2psZyMm.
2. S. Lohr, "For Big-Data Scientists, 'Janitor Work' Is Key Hurdle to Insights," *New York Times*, 17 Aug. 2014; nyti.ms/2kl1V3Y.
3. R. Delgado, "Why Your Data Scientist Isn't Being More Inventive," *Dataconomy*, 15 Mar. 2016; bit.ly/2rxoy73.

Seth Earley *is CEO of Earley Information Science (www.earley.com). He's an expert in knowledge processes, enterprise data architecture, and customer experience management strategies. His interests include customer experience analytics, knowledge management, structured and unstructured data systems and strategy, and machine learning. Contact him at seth@earley.com.*

Technology Policy and the Trump Administration

Shane Greenstein
Harvard Business School

Technology policy has been a low priority for most voters in presidential elections in the post-war era. The most recent contest was no exception. Arguments about technology policy never made it into campaign commercials, to say the least, nor even a minute of the televised presidential debates.

So it goes.

Many denizens of the high-tech world did not expect Donald Trump to win, woke up to his triumph, and suddenly wondered what impact his new policies might have on their business. Needless to say to anybody who paid attention, his campaign was not much help answering those queries, since he was not very specific about his technology policy.

This column considers two nonpartisan questions. How will his (likely) technology policies affect the value of US firms in information technology markets? Details about policy should become known in the first half of 2017, and they suggest a second question: what details should an investor care about, as the new administration hashes out the details?

Trade

Let's start with trade, which was a visible aspect of Trump's campaign. He expressed dissatisfaction with the position of the US in the world trade system. He focused on the exit of jobs in footloose manufacturing industries, the North American Free Trade Agreement (NAFTA), and China's mercantilist actions.

What can changes in trade policy do to IT firms? At a general level, every large US tech firm is integrated into the world framework for trade, so tearing up the system could cause considerable damage.

Generally speaking, every large US firm sources inputs from outside the US and sells final products outside the US. Every major software firm uses programmers in the US and outsources some amount of work to programmers outside the US, and, again, sells their products to buyers outside the US. That goes for Apple, Cisco, IBM, Microsoft, HP, Oracle, Google, Facebook, and on and on. These firms will lose market share and profits if the costs of inputs and labor increase, or if the number of potential buyers declines.

Since Trump's populist antitrade tirades conflict with his generally probusiness attitudes, we should expect quite a fight inside the administration over the practical details of trade policies. Whether input and labor costs will increase, and how much, and whether market buyers will decline, and how much, cannot be determined until those details get set.

As for Trump's dissatisfaction with Chinese mercantilism, most experts predict that his confrontational policy will go nowhere. I find these experts so persuasive that I'll take a bet—$100 to your favorite charity or mine. You win if a major US IT firm improves its mainland Chinese market share in the next four years.

Beyond that, one additional trade question is worth watching. Several US firms, including Cisco, IBM, Google, Apple, and Facebook, have a large foreign presence and would like to repatriate their foreign earnings as US dollars without paying taxes. The Obama administration refused to initiate such a tax holiday, and Trump might think differently. That would move stock prices if implemented.

That adds up to a big unknown for firm values. The value of virtually every US-based IT firm depends on the outcomes of these policy debates. That supports an approach for investors: expect volatility across the entire sector and adopt investment strategies to hedge against it.

Immigration

During the campaign, Trump focused hostility toward immigration. Attention was directed at unskilled immigrants from south of the US border, as well as those from nations with Muslim majorities. Quite frankly, if immigration from Mexico slows, then industries other than high technology—such as agriculture, service work, and construction—will be most affected. It is hard to see any direct effect for IT firms arising from that type of policy.

What if the US (deliberately) started processing visas from Muslim-majority countries with more laborious delays? It would affect the visits of foreign executives from countries such as Dubai, Saudi Arabia, and Malaysia. That might affect

 Published by the IEEE Computer Society

This article originally appeared in IEEE Micro, *vol. 37, no. 1, 2017.*

boutique tourist and consulting businesses, but that touches only the edge of US IT firms.

Investors should focus on policies for high-skill immigrants. Every major high-tech firm employs high-skill immigrants, and this group comprises founders for many venture firms. Many high-skill immigrants have master's degrees or PhDs from US universities. The present system works reasonably well for those with degrees.

Drastic changes—such as slowing the granting of visas and green cards to those with degrees—will slow down innovation in the commercial IT sector. That holds for virtually all US IT firms, so any slowdown hurts investments across the sector.

One aspect of high-skill immigration is difficult to forecast—namely, changes to the system for H-1B visas. The H-1B system is already rather constrained, and nobody expects that the Trump administration will try to reform it. More to the point, any tighter limits would hit a few firms hard. Will that happen? It is hard to say right now. Investors have to watch and wait.

Research and Development

A new administration can also change policy for R&D. Although it's less visible to the average voter, the US government is the single largest funder of basic science and also a large funder of experiments in applied science.

This matters to US IT firms, who have benefited from this funding in the past. For example, new network engineering, search engines, AI, and robotics can trace their invention to federal funding. In addition, many computer scientists got their first experiences on projects funded by this federal money, which effectively subsidized US technology workforce training.

There is no reason to expect the emergence of additional technologies to be any less sensitive to federal funding, so changes to the funding level at DARPA and the National Science Foundation are key budgets to watch. The same goes for R&D funding at the National Institutes of Health, NASA, and the Department of Energy, which also support R&D that works its way into commercial IT products.

That adds up to a straightforward forecast: if the budget for R&D at these agencies declines, that is bad for the whole sector. Drastic cuts slow down innovation, whereas growth speeds it up across the entire portfolio.

Commercial Policy

The administration can have immediate impact through staff and personnel appointments to agencies that make commercial policy for high tech. For example, most observers expect the administration to appoint directors to the Securities and Exchange Commission who oppose government actions. Expect government regulators not to stop Wall Street banks from returning to the kind of self-serving (and sometimes unethical) actions observed in the 1990s during the IPO boom. Frankly, I think this will be bad for the US startup economy.

Investors should also expect the Federal Communications Commission (FCC) to appoint decision makers who will not intervene in Internet markets. Net neutrality will not be enforced, and many other recent initiatives will be reversed, such as those aimed at opening up the set-top box for cable television.

That will raise the value of big cable firms, such as Comcast, and other carriers, such as Verizon, because it will give them the upper hand in negotiations on a range of issues, such as zero rating, interconnection fees for moving data into ISP networks, and collocation fees for content delivery networks. Again, frankly, I think this will be bad for the value of content firms with big data applications, such as Netflix, YouTube, Facebook, and venture capitalists backing new streaming entrants.

Antitrust is a more ambiguous area. The Obama administration let 99 percent of mergers go through, and that will continue. The only open questions concern big mergers.

Watch the early test cases for clues about the general approach of the new administration. On the campaign trail, Trump expressed dismay about the proposed merger of AT&T and Time Warner, but most of his circle of advisors are hostile to blocking mergers. So, the outcome to that merger proposal will show a lot. Another test case could be a proposed merger between T-Mobile and Sprint, which the Obama administration's appointees talked down before it was formally proposed. Will management attempt to resurrect it? We'll have to watch and wait.

Also watch the administration policy in privacy, which the FTC took a lead on in the last few years. The issues are varied, subtle, and difficult. The policy outcomes make an enormous difference to product design and the operations of many firms, especially those in online advertising and healthcare. Again, investors will have to wait and see.

Similarly, Apple faced a quandary protecting privacy when it negotiated with the FBI about breaking into an old iPhone. The FBI learned of another way into the phone, so the broad issues never got resolved. The Obama administration sided with law enforcement, and Trump did, too (even calling for a boycott of Apple). Expect more volatility from these issues. It is a wild card for values at many firms.

Overall, Trump's policies look like a mixed bag for the value of many US IT firms, and contain many dangers. This much I can forecast: most savvy tech firms woke up the day after the election and added staff to their Washington lobby organizations. The cynic in me also expects the Trump administration to try a quid pro quo, such as offering a tax holiday as a bribe to gain silence on other issues.

That suggests a somewhat partisan forecast. I do not expect that most CEOs want their issues to be invisible in the next election. I also expect their stance next time to depend on their experience in the next few years. MICRO

Shane Greenstein is a professor at the Harvard Business School. Contact him at sgreenstein@hbs.edu.

Cloud-Native Applications and Cloud Migration

The Good, the Bad, and the Points Between

Cloud-native features in your information technology (it) systems come with many advantages, but they also come with a cost. The costs vary greatly, based upon your applications and data. Sometimes being cloud native doesn't make economic sense, and sometimes it does.

Keep your eye on that ball as you migrate to the cloud.

A global 2000 company or a government agency typically has more than 5,000 applications on legacy platforms. These applications are placed into categories: ones that should move to the cloud and ones that should not. If they move to the cloud, then do they need to be altered to leverage cloud-native features (refactoring), or they could be moved with few or no changes (lift and shift)?

As you can see in Figure 1, an application can take many paths on its migration to the cloud. They're called the 7 R's of migration. This includes lift and shift (rehosting), or moving the applications with few or no changes. Partial refactoring means changing a small percentage of the code to leverage some cloud-native features. Complete refactoring means rewriting most of the applications so that they become cloud-native applications.

You need to pick a path from one of the 7 R's for each application, based upon its immediate, as well as long-term, return on investment. In some cases, you need to break applications apart into reusable components, and in other cases, applications can be replaced with a software as a service (SaaS) application analog. Some applications need to be removed, and many applications should stay put.

The toughest part of application migration to the cloud is figuring out the correct path. The retire path is usually obvious from the start. If applications are movable to the cloud, how should they be configured, changed, or not? Enterprises still struggle with these issues, but best practices are starting to emerge.

Why Go Native

The pros of going to cloud-native features include the following:

- **Performance.** You'll typically be able to access the native features of public cloud services to

EDITOR
DAVID S. LINTHICUM
david@davidlinthicum.com

 Published by the IEEE Computer Society 2469-7087/18/$33.00 © 2018 IEEE

provide better performance than is possible with nonnative features. For example, you can deal with an input/output (I/O) system that works with autoscaling and load-balancing features.

- **Efficiency.** Cloud-native applications' use of cloud-native features and application programming interfaces (APIs) should provide more efficient use of underlying resources. That translates to better performance and/or lower operating costs.
- **Cost.** Applications that are more efficient typically cost less to run. Cloud providers send you a monthly bill based upon the amount of resources consumed, so if you can do more with less, you save on dollars spent.
- **Scalability.** Because you write the applications to the native cloud interfaces, you have direct access to the autoscaling and load-balancing features of the cloud platform.

The price you pay for these advantages is portability. Applications that are localized for specific cloud platforms are not easily ported to other cloud platforms. Doing so takes a great deal of rewriting or refactoring of the code. For all practical purposes, you are locked into that cloud platform.

If applications run on the target cloud platform for years, you're bound to get your investment back in code changes and testing. You have to look at the advantages of going cloud native case by case and application by application.

Unfortunately, there are no premade checklists you can use to make these assessments. You simply have to get as smart as you can about the trade-offs and make the best decisions you can (https://www.cloudtp.com/doppler/pros-cons-going-cloud-native/).

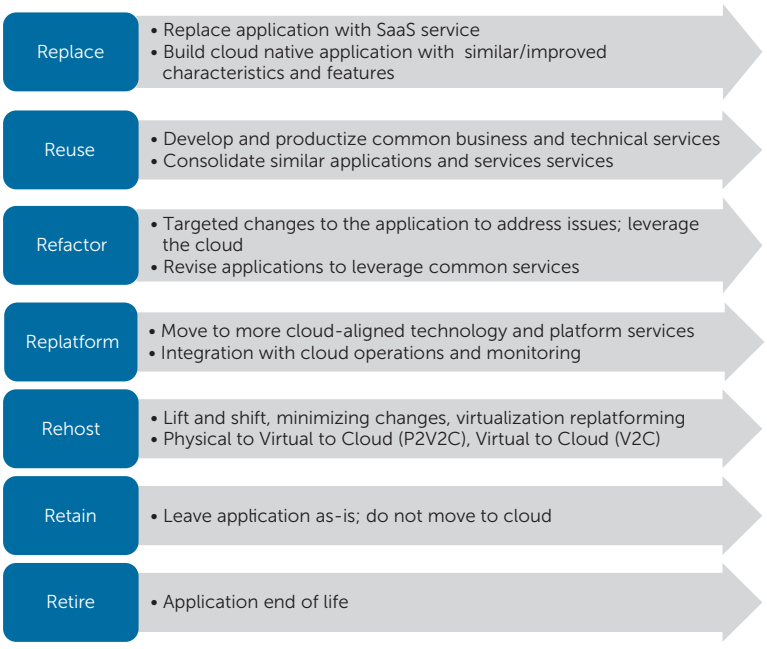

FIGURE 1. The 7 R's of cloud migration. Source: Cloud Technology Partners

A General Approach to Becoming Cloud Native

To take proper advantage of a cloud platform, including infrastructure as a service and platform as a service (PaaS), you have to design the applications so that they're decoupled from any specific physical resource. For example, if you access I/O directly from a platform as Linux, you need to access the cloud's abstraction layer, or its native APIs.

Clouds can provide an abstraction or virtualization layer between the application and the underlying physical (or virtual) resources, whether they're designed for cloud or not. But that's not good enough. If you're truly going cloud native, you need to directly control the native resources.

When this architecture is considered in the design, development, and deployment of an application, the utilization of the underlying cloud resources can be as much as 70 percent more efficient. This cloud computing efficiency equals money. You're paying for the resources you use, so applications that work more efficiently with those resources run faster and generate smaller cloud service bills at the end of the month. Several cloud-cost governance tools can monitor these costs for you. I highly recommend using one or more of these tools if you plan to spend the time and money to move and refactor applications to be cloud native.

The Downsides of Cloud Native

There are trade-offs to cloud-native application development and design. Be aware of the trade-offs before you bind your application to a specific cloud. You should do this with the heart of a chief financial officer, not a technologist. At the end of the day, this is about costs, not what is technically superior.

The biggest trade-off, as we covered already, is that you're giving up portability for the advantages of being cloud native. Applications that are localized or made native for specific cloud platforms are not portable to other cloud

This article originally appeared in IEEE Cloud Computing, *vol. 4, no. 5, 2017.*

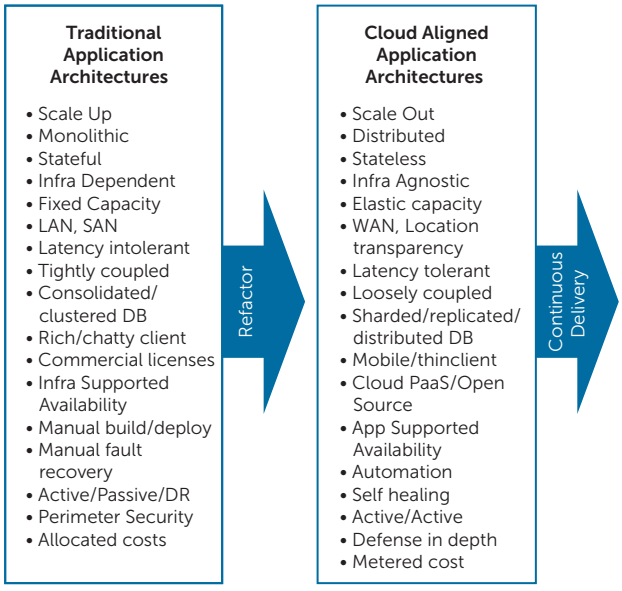

FIGURE 2. Being cloud native means more than changing code. You need to change your application architecture as well. Source: Cloud Technology Partners.

platforms without a great deal of rewriting or refactoring of the code.

That being said, few cloud-to-cloud migrations have happened yet. This is true even when rehosting or lift-and-shift migration paths are chosen. If the current trends continue, the argument against lock-in does not seem to be much of an issue for most enterprises. But you should at least keep this issue in the back of your mind.

Cloud lock-in does expose the enterprise to some risk if the cloud platform provider becomes difficult to deal with, raises prices, or most likely, some of the services you've purchased and use get turned off. In those cases, you're going to have to bite the bullet and re-refactor those applications to be cloud native on another cloud.

Some developers are getting good at placing the cloud-native features of an application into a specific domain of the application design. This allows them to more easily change the application for other cloud platforms, if needed. These good application design practices are not common; most developers exploit the cloud-native interfaces systemic to the application. Those are harder to change.

My best recommendation is that you at least consider cloud-native application development approaches and best practices—if not for the efficiency, then perhaps for cost or performance. Although there is a clear trade-off in portability, the benefits seem to outweigh that cost quickly if you're going to run the application for more than a couple of years.

Emerging Architectures

Cloud native is not about just changing the code to follow the features of a specific cloud; it's also about changing your approach to architecture design. As you can see in Figure 2, what's emerging is a world of cloud-aligned application architectures.

These cloud-aligned architectures can autoscale and are distributed, stateless, and loosely coupled, to name a few features. If you want to make applications truly cloud native, the architecture must be rethought before you begin refactoring the code.

Does this new approach to application architecture suck resources and money, as well as adding a great deal of risk? Yes. However, the risk-reward scale typically leans to the reward side if the life of an application is 10–15 years (which it is for most enterprises). The effort of both rearchitecture and refactoring for an application with long-term use will pay for itself many times over.

However, enterprises in the US are not wired for long-term investiture. Most enterprises opt for lift and shift versus refactoring for cloud native. Enterprises are rewarded based upon earnings, and the lower the cost of new software development and IT, the more earnings they can claim—short term.

Do the Right Thing

The evidence is compelling for the cloud-native application path when migrating applications to the cloud. The benefits outweigh the costs for most applications that are picked to be moved to the cloud, but given that refactoring costs 30 times simple rehosting, enterprises are reluctant to jump in with both feet.

So, this will be another learning process. Much like we saw when we made applications platform native, such as Win32- and Unix/Linux-native APIs, we had to fail first. In other words, working around the native platform features led to applications that did not live up to the expectations of the business, and IT had to go back to square one to redesign and refactor the applications to meet the needs of the business.

I suspect we'll follow the same patterns here. In a few years, cloud native will become the best practice. However, that won't happen until we fall on our faces a few more times. Some things never change.

Dana Ulery: Pioneer of Statistical Computing and Architect of Large, Complex Systems

Irina Nikivincze
Georgia Institute of Technology

From Grinnel to JPL

Dana Lynn Ulery (1938–present) did not grow up in a typical 1950s family.[1] Her mother, Meriam Mueller[2] (1908–2005), was a businesswoman and active volunteer in the local hospital. Her father, Harry Tanzer, died when she was only 2 years old. Over the years the loneliness that Dana felt was replaced by her love for learning and school. Graduating from Grinnel, a small liberal arts college in Iowa, in 1959 as a double major in mathematics and English literature, Dana's future looked uncertain, but exciting. She was getting married and the young family was going to tour the West and move to Pasadena, California.[3] As with many other women of that time her choices were limited—she could be a secretary, nurse, stewardess, or teacher. Just in case, she got a teaching certificate and quite unexpectedly got an offer to teach mathematics in a high school in Burbank, California, minutes away from her husband's workplace. Her initial excitement soon dissipated after she learned that the class that she was teaching was filled with kids who could not read. There was more to it. All of those kids were boys from a reform school—teenagers who already had gotten in trouble with the law. A young female math teacher with her rules and homework was the least of their concerns and perhaps a source for amusement. The boys refused to do homework and Dana's teaching aspirations vanished day by day as her focus shifted from content to class discipline. One year was enough for Dana to realize that teaching was not what she wanted to do in her life and that she needed a new job.

In Pasadena, her husband, Harris Ulery, just started his graduate studies in organic chemistry at the California Institute of Technology. Since Caltech was a technical school, Dana was convinced that somewhere there was a job for a math major. Having gotten her confidence together, Dana walked the halls of Caltech, stopping by open doors and inquiring about jobs. "Somebody finally took pity at me and sent me to JPL," she later recalled.[4]

The National Aeronautics and Space Administration (NASA) Jet Propulsion Laboratory (JPL), a large federally funded research center, was located near Caltech. Even though JPL hired female mathematicians in the 1940s to 1960s to perform trajectory calculations, it was in an all-female unit.[5] There was no precedent for hiring women for an engineering job. A man at JPL who interviewed Dana was amused by an enthusiastic and math-minded young lady. Nevertheless, he offered her a job and took care of the paperwork, so she did not need to go to Human Resources. Dana became the first and only female engineer at JPL in 1960. JPL did not hire other women engineers for the next 7 or more years. One of Dana's first assignments was to evaluate the systematic pointing errors in the Goldstone Polar-Mount Antenna using the star-track data.[6] During the following three years she worked on real-time tracking systems and algorithms for NASA's Deep Space Network—the work that was exciting and that she became very fond of.

The official title that Dana held was of a junior engineer and it meant that she was paid less than her colleagues. Having learned that from a friend, Dana succeeded at changing her title (and the corresponding pay) to one of a Research Engineer. She greatly enjoyed her job at JPL but not for long. In 1963, Harris was graduating with his PhD and announced that he would be looking for jobs. Although, Dana wanted to stay at Pasadena, the family weighted in, and the move was imminent.

Among the three job offers that Harris received, the least "horrible" in Dana's opinion, was one in Delaware, OH. Dana stayed at JPL as long as she could and then it was time to move to the East Coast. Moving to a new place meant facing the same obstacle—with no established networks, getting a new technical job would be a hurdle. She explored every opportunity: a job at Thiokol, a chemical company that also contracted with NASA, and a programming job at Getty Oil. Her training and previous experience at JPL made her a superior candidate for the programming job. Again, she became the only woman in her team. Eight men who worked in her group all had desks, but since they were not used to women and did not know how to treat her, Dana was given a private office. The unfortunate side of this arrangement was that she became isolated. There were no mental challenges, contacts, or discussions of

2469-7087/18/$33.00 © 2018 IEEE Published by the IEEE Computer Society

interesting problems. The new job did not come close to the work she was doing at JPL. The conflict at home was growing: Dana wanted to go back to school, while Harris wanted his life to continue as it was without any changes. Dana pursuing graduate work would put a burden on the whole family, and especially on him. The family somehow managed when he was a graduate student and Dana worked, took care of the house, and their small son. However, Harris was not ready to assume some of Dana's responsibilities. Divorce was imminent and only the birth of their second child, Terrie in March 1965, made Dana halt her aspirations.

Going back to graduate school at the University of Delaware

Five years later, in 1970s, at the age of 32, Dana entered the graduate program at the University of Delaware. She felt revived to be back in an intellectual environment. It meant a great deal to her to be around people who talked about interesting problems. One of them was a young faculty member Dr. Hatem Khalil. His openness and confidence would encourage Dana to stay in the program beyond Master's.[7] Hatem Khalil defended his dissertation the same year Dana entered university under Dr. John H. Giese, Professor of Mathematics and Computer Science and Chief of Computing Laboratory at Aberdeen Proving Grounds, Maryland. Khalil was a mathematician who was interested in computers and the new emerging field of computer science. He would become not only a teacher but also a colleague and friend.

Khalil also was one of most difficult professors in numerical analysis. In the early 1970s, there were very few female students in the department. Men would talk a lot in classes, and they appeared smart and knowledgeable. Dana got so intimidated that she took it to her advisor. She thought that she was doing poorly as she could not think and follow the class discussion as fast, and she was pondering if she should even stay in this class. To her surprise, Khalil told her that she was the best student in his class and even though others were talking a lot and answering his questions, they were giving wrong answers. These comments became a source of relief and encouragement. In the company of Dr. Khalil, Dr. Giese (Dana's co-advisor), and her husband Dr. Harris Ulery, the work became fun. "If I would not have these men, I would not have stayed in the field," she later recalled. At home, Dana still took care of the family and cooked dinners, and the family ate together. Not having a full-time mother meant that the family had to cope with late dinners, accommodations for kids, and other "inconveniences." Only with their support and cooperation was Dana able to get through school.

In 1975, Dana had finished her dissertation "Computer Science's Reincarnation of Finite Differences," submitted in 1976 for the degree of Doctor of Philosophy in Applied Science (Computer Science)." In 1976, her advisor was doing lecturing and research at the American University in Cairo in Egypt. Not having other prospects, Dana accepted a postdoc position at Cairo University and became a visiting lecturer. At that time, the subject of computer science intrigued academics around the world, including in Egypt. However, that temporary position did not resolve the question of what she should do next.

Dana's graduate training was grounded in mathematics, statistics, and computing. Her graduate department trained students in "Applied Sciences," although it was already called the department of Statistics and Computer Science.[8] Dana's advisor, Dr. Khalil, was involved in curriculum reform at his department. Unlike some of his colleagues, he gave importance to the theory and practice of computing. In his view, programming was to computer science "what the laboratory is to the physical sciences," and the study of computer science that he defined as "the theory and practice of programming computers" should include the study of algorithms—the main theoretical foundation for programming.[9] Thus, Dana's graduate work was in the theory of programming—in algorithms. As part of her Master's degree she devised a language-oriented system that she called Lineal for solving problems in linear algebra. In her dissertation she explored the techniques for solving partial differential equations, in particular the method of finite differences. The repetitive nature of this method begged for a generalization and automation. That was exactly what Khalil and Giese were doing—devising and exploring algorithms for generating families of difference approximations. In her work, Dana extended their method to higher dimensions and explored its boundary conditions.[10]

What to do with her PhD still remained a question. As with many other women, Dana did not get much direction from the

university about what to do next. Universities and organizations for the most part were ambivalent towards women.[11]

DuPont, 1977–1994

One day somebody called the department at the University of Delaware looking for students interested in high-level languages and contract work. The department gave Dana's name. Don (Donald) Marquardt was a manager at DuPont who was highly respected and influential in statistics. Marquardt told Dana about the job, but she did not find it very interesting. When Dana's advisor found out that she refused the job, he was furious. Fortunately, DuPont called back and this time they had a more interesting offer that Dana enthusiastically accepted.

Dana was working in the Applied Statistics Group of the Engineering Department of "E. I. du Pont de Nemours and Company" in Wilmington, DE first as a software engineer and later as a consultant. DuPont was a large high-technology company that also had "large and conceptually complex problems."[12] It was one of the best places for basic research. The company was collecting the Crème-de-la-Crème—technical people from various places. Very prestigious PhD statisticians worked there. Donald Marquardt managed a group of approximately 30 people, of whom 18 were statisticians, half had PhDs, and the rest had MS degrees.[13]

The Applied Statistics Group of the Engineering Department wanted to do more work with computers. Working in industry, Marquardt wanted statisticians to be more entrepreneurial and business-oriented. His unit was providing "statistical consulting" services at DuPont and serving close to 110 sites in the United States. He encouraged consultants to get involved in projects, learn the lingo, but avoid doing the engineering work or project management, and instead address a client's problem using statistical methods. Frequently, the end product of statistician's work became "embodied in a computer program that the client could use to handle future instances of the same problem."[14] A computer program was not only able to standardize and monitor manufacturing processes, but also improve the organization of operations. Maintaining some distance from actual projects put consultants in an advantageous position to be able to codify and improve the organizational processes. In addition to being quality control designers, the role of statistical consultants became one of system builders and organizational optimizers.

By the time Dana entered DuPont the company already had problems. The business was in trouble because the quality of the products was inconsistent. A suggestion was made to put together a large system that could help to get the product quality back in order. Since 1972, DuPont was developing a system for product quality management that incorporated business philosophy, management, and technology systems. The development of a product quality control system by Marquardt's team was an important project that required innovations in statistical methods and its implementation as a computer program. If in a pre-computer age manufacturers could use Shewhart charts to monitor process changes, but with increase in scale and complexity of production, it was difficult to discover intermittent changes. Marquardt's team proposed a cumulative control strategy—one that they called Cumulative Sum (CUSUM), which included additional measures, such as warning limits, run tests, and more sophistical detection of changes through average run length curves.[15] The development and use of such a complex quality control system depended on the use of computers and on the training of the personnel.

The Applied Statistics Group was running simulations and modeling, but they did not have many people who could program it on PDP-11. Dana worked on computer implementation of CUSUM quality control schemes that were able to detect changes in processes ahead of production defects.[16] By 1977 that project became the major consulting activity of Marquardt's group: they trained over 15,000 people in product quality management and implemented over 10,000 computerized CUSUM control loops.[17] It also launched Dana's career in software quality management.

In the Research Division DuPont had very specialized personnel—some of whom had PhDs specifically in color theory. Although they collected data and did experiments, they did not solve nonlinear equations that were Dana's area of expertise. However, working with them did not get her closer to statistical work and the opportunity to contribute to that area. Instead, she moved to software engineering because it was "available" to her. Together with new career advancement came changes in family. Old arrangements were

not working. Dana divorced her husband in 1978 and married her colleague William Fellner, a consultant with DuPont's Quality Management and Technology Center in 1980, with whom she shared many interests and hobbies.

Dana excelled in development of large systems and applying new tools for the optimization of existing systems. She wrote a quality system named QFACS that facilitated data communication. As a result, product pigment quality improved, and she was asked to do a major presentation. As the first quality system was gradually aging, DuPont wanted a new quality system, only much bigger. That system included close to 150,000 people. She was working in the team with two other men and was responsible for software architecture of the system. At that time they did not have software architects but they had Dana. Such system had to handle not only operational complexity but also coordinate an influx of data from different business units. In 1991, together with Donald Marquardt, Dana edited a 600-page reference book entitled *Product Quality Management.*[18]

In the mid-1990s, DuPont, which was made up of many businesses, started to outsource many of its parts to other companies, such as Computer Sciences Corporation (CSC) and Anderson Consulting. They kept a small number of people on a contract basis. The outsourcing fever was only starting, but it was clear that sooner or later this move would significantly reduce DuPont's internal consulting unit. Dana lost her job in 1994. At that time she was 56 years old. Even with her experience, finding another job was not easy.

United States Army Research Lab (ARL)

After many phone calls and networking, including with former advisors, she managed to get a job at a senior position at the Aberdeen Proving Ground, MD. She became a senior research scientist and Acting Chief of Intelligent Systems Branch at the United States Army Research Lab. As in DuPont, she continued doing applied research, except now her focus shifted to systems that helped the United States military to do things better and faster: the use and impact of intranets,[19] software prototype for purchasing (BuyIt)—a part of the Corporate Business Application Software System (C-BASS), and finally, the technologies for intelligent data/knowledge fusion that facilitated actionable knowledge.[20] Dana approached system design with a lot of consideration and responsibility because the systems that she was designing were shaping users and their thinking.

Having worked on many complex projects, Dana learned that technical and social aspects of software engineering were often intertwined. Problems appeared to be technical only at the beginning and eventually came down to dealing with people. Such systems often required streamlining existing work processes while organizations resisted changes in business practices. Even the small-scale transactional systems that she worked on required improvements in work practices. The management of work practices necessitated the management of organizational culture. By that point, Dana became increasingly interested in social aspects of systems, growing semantic web (internet) and policy issues. From the late 1980s, Dana worked on common data standards while serving on the Accredited Standards Committee X12 (ASC X12) of the American National Standards Institute (ANSI) and as Pan American Delegate to the United Nations Electronic Data Interchange for Administration, Commerce, and Trade (UN/EDIFACT). In 2007, she retired from her position at the United States ARL.

Conclusion

As many other pioneering women, Dana had few expectations and little knowledge of career options in technical fields. Although she overcame some gender stereotypes by entering a technical field and by working on technical problems, she could not change the environment or the behavior of other people. Institutions treated women differently: they often occupied separate spaces, were not given client accounts and money that their male colleagues had, and suffered the indignity of having men even refuse to work with them. It was easy to be in "someone's way," and if a woman moved up to a position of responsibility—somebody did not like it. Men's attitudes often were very wearing on technical women. If things were different, Dana thought that she might have gone higher, for example, became a manager or a director of a larger part of the organization—a dream that she partially realized at the ARL.

The biography of Dana Ulery provided a glance at the transformational role of computers in science, industry and military at the end of the 20th century. It revealed that computers were indispensable in solving complex quality control issues in the 1970s,

optimizing both technical and organizational processes and, perhaps, facilitating organizational outsourcing in 1980s. Dana's unique skills, which combined the knowledge of statistics, programming, and system architecture, came in handy in addressing those problems. Dana had a successful career in industry. She followed the opportunities. One technical challenge led to another. Her career allowed her to see the breadth of technology much more than she would have encountered in a classroom. She was able to raise two children and finish graduate school. Even long hours at graduate school had a positive side effect—it allowed her kids to see her as a full human being. Her daughter later became a scientist with a PhD in Biology. Being one of the first in computer science, Dana Ulery managed to have what other women of that time could only dream of—a loving family and the excitement of intellectual work.

References and Notes

1. This biography is based on the interview with Dana Ulery conducted by the author on December 4, 2013 for the project "Careers and Contributions of the First Doctoral Women in Computer Science" sponsored by the ACM History Committee.
2. "Meriam Mueller," *The News Journal* (Wilmington, DE), November 7, 2005, p. 13.
3. "Dana Tanzer Married to Harris E. Ulery," *St. Louis Post-Dispatch*, August 30, 1959, p. 102.
4. D.L. Ulery, "Making Science: Careers and Contributions of the First Doctoral Women in Computer Science," Interview by Irina Nikivincze, Landenberg, PA, December 4, 2013.
5. E. Conway, "Women Made Early Inroads at JPL," March 27, 2007, https://www.jpl.nasa.gov/news/news.php?feature=1327.
6. D. Ulery and J. Fearey, "Evaluation of Goldstone Polar-Mount Antenna Systematic Errors from Star Tracks," Technical Memorandum 33-45 (Unclassified), Jet Propulsion Laboratory, May 5, 1961, *Publications of the Jet Propulsion Laboratory July 1961 through June 1962 (Bibliography No. 39-3)*, S.L. Kresser and R.J. Sippel, eds., Jet Propulsion Laboratory, October 15, 1962, p. 47.
7. D.L. Ulery, "Computer Science's Reincarnation of Finite Differences," PhD dissertation, Univ. of Delaware, 1976.
8. H.M. Khalil, "A Two-parameter Family Of Approximations to The Two-dimensional Heat Equation," PhD dissertation, Univ. of Delaware, 1970.
9. H. Khalil and L.S. Levy, "The Academic Image of Computer Science," *ACM SIGCSE Bulletin* 10, no. 2, 1978, pp. 31–33.
10. See Ulery's dissertation and J.H. Giese, H.M. Khalil, and D.L. Ulery, "Multiparameter Families of Differences Approximations for the First Initial Boundary Value Problem for the Heat Equation in an Arbitrary Region," *J. Eng. Math.*, vol. 12, no. 2, 1978, pp. 97–114.
11. This comment was made more than once in interviews conducted by the author with women pursuing degrees in computer science in early 1970s.
12. D.W. Marquardt, "Statistical Consulting in Industry," *Am. Statistician*, vol. 33, no. 3, 1979, p. 105.
13. D.W. Marquardt, "Statistical Consulting in Industry," *Am. Statistician*, vol. 33, no. 3, 1979, pp. 102–107.
14. Marquardt, "Statistical Consulting in Industry," p. 103.
15. D.W. Marquardt, "New Technical and Educational Directions for Managing Product Quality," *Am. Statistician*, vol. 38, no. 1, 1984, pp. 8–14.
16. D.L. Ulery, "Software Requirements for Statistical Quality Control," *IFAC Real Time Programming*, 1985, pp. 39–42.
17. Marquardt, "New Technical and Educational Directions for Managing Product Quality."
18. D.W. Marquardt and D.L. Ulery. *Product Quality Management*, E.I. du Pont de Nemours, Quality Management & Technology Center, 1991.
19. D.L. Ulery, *ARL Intranet Analysis and Development Study*, ARL-CR-441, Army Research Laboratory/Georgia Institute of Technology, 1999, http://handle.dtic.mil/100.2/ADA362869.
20. R. Scherl and D. L. Ulery, *Technologies for Army Knowledge Fusion*, Army Research Laboratory, 2004, http://purl.access.gpo.gov/GPO/LPS125004.

Irina Nikivincze *is a Postdoctoral Researcher at the Georgia Institute of Technology. Her research explores scientific careers, gender, achievement, and recognition in computer science. Contact her at irina.nikivincze@amac.gatech.edu.*

This article originally appeared in IEEE Annals of the History of Computing, *vol. 39, no. 2, 2017.*

by Charles Day

Computing in World War I

This year marks the 100th anniversary of America's entry into World War I. On 6 April 1917, both chambers of Congress voted overwhelmingly to grant President Woodrow Wilson's request to declare war on the German Empire. Among the first US military formations to arrive in Europe was US Battleship Division Nine, which comprised the battleships *Delaware, Florida, New York*, and *Wyoming* and the destroyer *Manley*. The ships joined the Royal Navy's Grand Fleet at its base in Orkney on 9 December.

The fastest, most modern warships of the time were propelled by oil-powered steam turbines. Short on oil because German U-boats were sinking British oil tankers, the Royal Navy had asked the US Navy to send older, coal-powered ships instead. Despite their relative age, the battleships of Division Nine were formidable war machines. All of them survived the war.

When the US ships went on training exercises with British ships, it became clear that the British could sustain higher and more accurate rates of fire. Part of the British superiority arose from the three years of experience the Royal Navy had already gained from fighting Germany and its allies at sea. But the Royal Navy also had a technological advantage: the main guns of its largest battleships were controlled by a computer.

Computer control was an integral feature of a new type of battleship introduced a decade before war broke out. Named after their archetype, HMS *Dreadnought*, the new battleships were equipped with ten or so long-range guns of a single large caliber, rather than a mix of small, medium, and large calibers. If a target's range, bearing, and speed were known, an aiming command could be transmitted to the gun turrets to direct a single, devastating salvo at the target. And if the aim was off target, a single adjustment sufficed for the second salvo.

The sequence of events that culminated in a *Dreadnought* salvo began with the use of an optical rangefinder, a binocular-like device that used triangulation to determine the range. Measurements from the rangefinder, including speed and bearing, were fed into a mechanical computer and electrically transmitted to mechanical computers that used lookup tables to determine the guns' elevation and azimuth at the moment of firing.

The biggest naval battle of World War I, Jutland, took place in the North Sea between 31 May and 1 June 1916. One hundred and fifty-one warships of the Grand Fleet engaged 99 warships of the Imperial German Navy's High Seas Fleet. Despite the Grand Fleet's numerical advantage and despite the High Seas Fleet's lack of computerized fire control, Britain lost more ships: 14 versus 6. Still, the German losses were deemed so heavy that the High Seas Fleet remained at port for the rest of the war. The Royal Navy continued to blockade Germany, while the Imperial German Navy resorted to unrestricted submarine warfare, a strategy that contributed to the US decision to enter the war on the side of Britain and her allies.

One of the British battleships that fought at Jutland was HMS *Warspite*. The ship not only survived the battle, it went on to serve in World War II with upgraded weaponry and a faster, though still analog, computer control system. During the Battle of Calabria on 9 July 1940, the *Warspite*'s fire control system succeeded in hitting the Royal Italian Navy's battleship *Giulio Cesare* at a range of approximately 24 km. The feat remains one of the longest-range hits in the history of naval gunnery. ■

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. 19, no. 4, 2017.*

 Published by the IEEE Computer Society 2469-7087/18/$33.00

COMPSAC 2018

Tokyo, Japan

July 23-27

Staying Smarter in a Smartening World

IEEE *Advancing Technology for Humanity*

IEEE computer society

Call for Papers

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 2018 is Staying Smarter in a Smartening World.

Computer technologies are producing profound changes in society. Emerging developments in areas such as Deep Learning, supported by increasingly powerful and increasingly miniaturized hardware, are beginning to be deployed in architectures, systems, and applications that are redefining the relationships between humans and technology. As this happens, humans are relinquishing their roles as masters of technology to partnerships wherein autonomous, computer-driven devices become our assistants. What are the technologies enabling these changes? How far can these partnerships go? What will be our future as we deploy more and more "'things" on the Internet of Things – to create smart cities, smart vehicles, smart hospitals, smart homes, smart clothes, etc.? Will humans simply become IoT devices in these scenarios and if so, what will be the social, cultural, and economic challenges arising from these developments? What are the technical challenges to making this all happen – for example, in terms of technologies such as Big Data, Cloud, Fog, Edge Computing, mobile computing, and pervasive computing in general? What will be the role of the 'user' as the 21st Century moves along?

COMPSAC 2018 is organized as a tightly integrated union of symposia, each of which will focus on technical aspects related to the "smart" theme of the conference. 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. A highlight of the conference will be plenary and specialized panels that will address the technical challenges facing technologists who are developing and deploying these smart 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. All submissions must adhere to IEEE Publishing Policies, and all will be vetted through the IEEE CrossCheck Portal.

Standing Committee Chair: Sorel Reisman, California State University, USA
Steering Committee Chair: Sheikh Iqbal Ahamed, Marquette University, USA

General Chairs: Shinichi Honiden (NII, Japan)
Roger U. Fujii, Fujii Systems, 2016 IEEE Computer Society Prelident

Program Chairs in Chief:
Jiannong Cao (Hong Kong Polytechnic University, Hong Kong)
Stelvio Cimato (University of MIlan, Italy)
Yasuo Okabe (Kyoto University, Japan)
Sahra Sedighsarvestani (Missouri University of Science & Technology, USA)

Workshop Chairs: Kenichi Yoshida (University of Tskuba, Japan)
Ji-Jiang Yang (Tsinghua University, China)
Hong Va Leong (Hong Kong Polytechnic University, Hong Kong)
Chung Horng Lung (Carleton University, Canada)

Local Chair: Hironori Washizaki (Waseda University, Japan)

Important Dates

Workshop proposals
Due date: 15 October 2017
Notification: 15 November 2017

Main Conference papers
Due date: 15 January 2018
Notification: 31 March 2018

Workshop papers
Due date: 10 April 2018
Notification: 1 May 2018

Camera Ready and Registration
Due date: May 15, 2018

PREPARE TO CONNECT

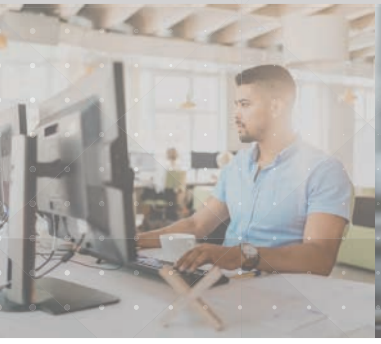
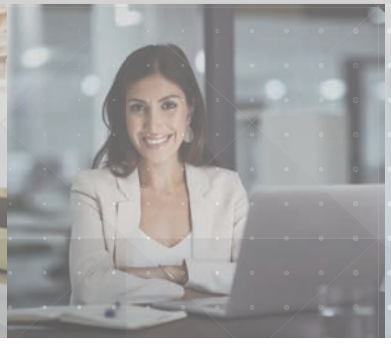
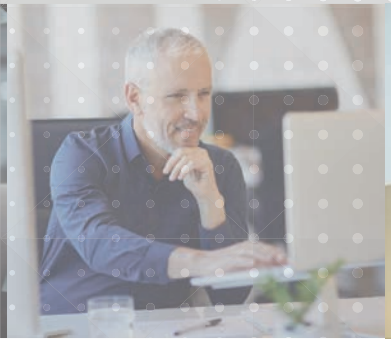


The IEEE Computer Society is launching **INTERFACE**, a new communication tool to help members engage, collaborate and stay current on CS activities. Use **INTERFACE** to learn about member accomplishments and find out how your peers are changing the world with technology.

We're putting our professional section and student branch chapters in the spotlight, sharing their recent activities and giving leaders a window into how chapters around the globe meet member expectations. Plus, **INTERFACE** will keep you informed on CS activities so you never miss a meeting, career development opportunity or important industry update.

Launching this spring. Watch your email for its debut.

IEEE computer society

INTERFACE