COMPUTING

edge

> Healthcare
> Cloud Computing
> Graphics and Visualization
> Robotics

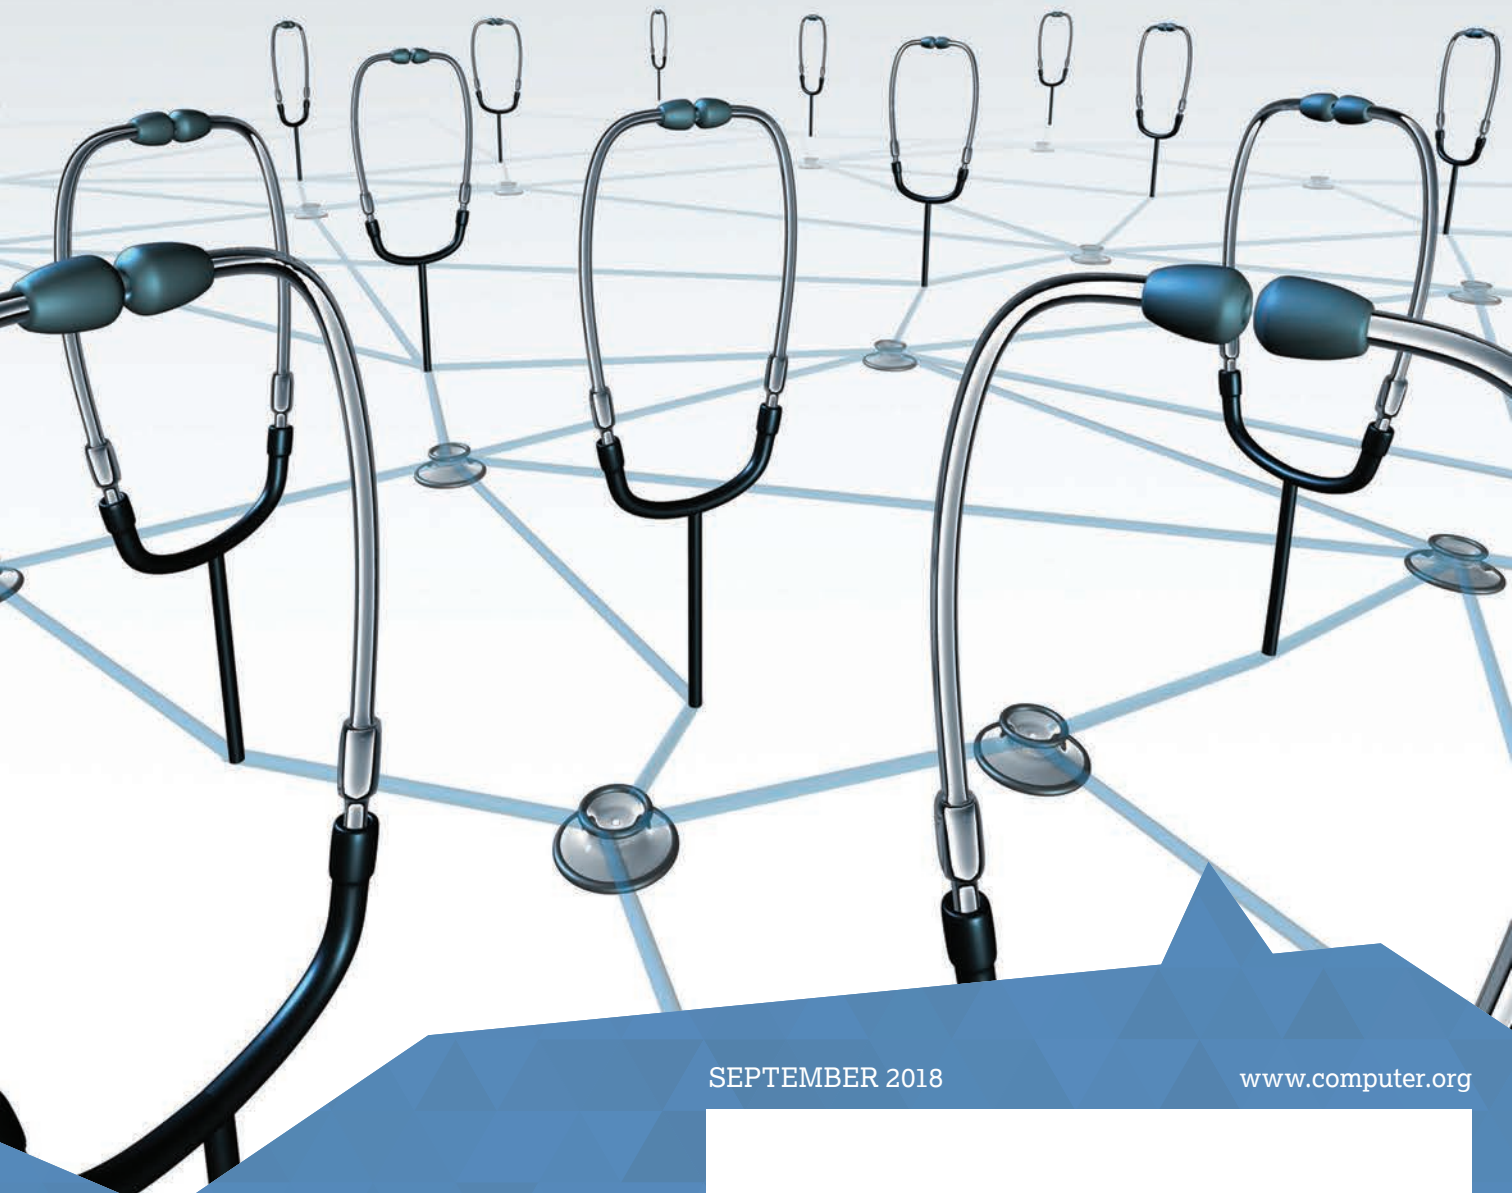

SEPTEMBER 2018 www.computer.org

IEEE
IEEE computer society

IEEE

LETTERS OF THE COMPUTER SOCIETY

IEEE Letters of the Computer Society (LOCS) is a rigorously peer-reviewed forum for rapid publication of brief articles describing high-impact results in all areas of interest to the IEEE Computer Society.

Topics include, but are not limited to:

- software engineering and design;
- information technology;
- software for IoT, embedded, and cyberphysical systems;
- cybersecurity and secure computing;
- autonomous systems;
- machine intelligence;
- parallel and distributed software and algorithms;
- programming environments and languages;
- computer graphics and visualization;
- services computing;
- databases and data-intensive computing;
- cloud computing and enterprise systems;
- hardware and software test technology.

LOCS offers open access options for authors. Learn more about IEEE open access publishing:

www.ieee.org/open-access

Learn more about LOCS, submit your paper, or become a subscriber today:
www.computer.org/locs

COMPUTING edge

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

Printed with inks containing soy and/or vegetable oils

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, *Drexel University*

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, *University of Vienna*

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

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

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

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*

SEPTEMBER 2018 • VOLUME 4, NUMBER 9

COMPUTING edge

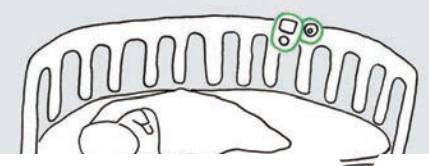

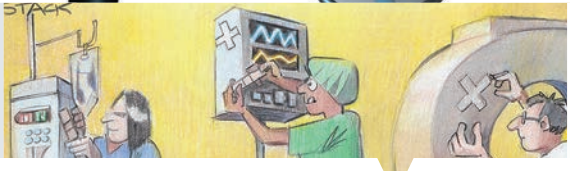

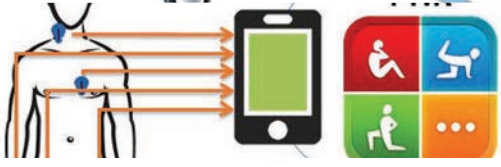

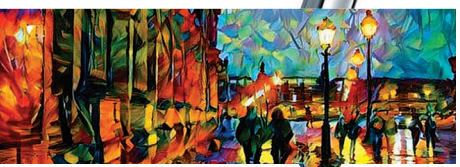

33

Application of Machine Learning to Computer Graphics

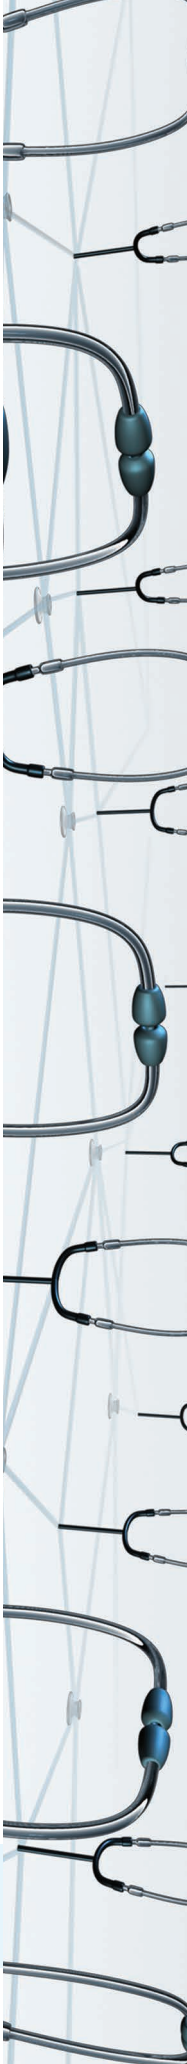

Magazine Roundup

Editor: 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

Attribute-Based Credentials for Privacy-Aware Smart Health Services in IoT-Based Smart Cities

Smart city–based IoT devices enable the collection of vast amounts of data, which can be used to provide more efficient public and private services. Among these, healthcare is especially relevant, and smart health (s-health) models are already being deployed. The authors of this article from the July 2018 issue of *Computer* propose attribute-based credentials (ABCs) to cope with s-health privacy issues and to set the stage for further adoption in other privacy-aware IoT-based smart city services.

Computing in Science & Engineering

Leveraging Cloud Computing for In-Silico Drug Design Using the Quantum Molecular Design (QMD) Framework

The authors of this article from *Computing in Science & Engineering*'s July/August 2018 issue present quantum molecular design, a novel cost-saving automated framework for de novo computational drug design. This technology not

 Published by the IEEE Computer Society

only addresses many of the challenges faced in the computer-aided drug design field by using highly accurate physics-based models, it also dramatically lowers costs by leveraging an AI heuristic search algorithm with targeted chemical space.

IEEE Annals of the History of Computing

A Slice of Norway's Computing History

Yngvar Lundh was instrumental in the development of some of Norway's earliest digital computers, having started his investigation of digital electronics in the 1950s. He continued his involvement with digital computers and digital communications for the next 40 years. The pioneering digital work at Norwegian Defence Research Establishment is the focus of Lundh's anecdote, particularly the hardware development activities and then early involvement with internetworking. Read more in the April–June 2018 issue of *IEEE Annals of the History of Computing*.

IEEE Cloud Computing

CUP: A Formalism for Expressing Cloud Usage Patterns for Experts and Non-Experts

The proliferation of cloud services, from infrastructure servers to software, has led to patterns in service deployment and provisioning practices, but not to standards for expressing and communicating such patterns to a broad-based audience. To enable the formal description and pattern classification of scenarios where cloud services are combined and provisioned to end users, the authors of this article from the May/June 2018 issue of *IEEE Cloud Computing* propose a textual and visual formalism: the Cloud Usage Patterns (CUP) formalism. With CUP, both general users and cloud experts can express patterns. By expressing patterns seen in practice, the authors demonstrate that CUP is practical and eliminates lengthy prose descriptions that result in misunderstandings.

IEEE Computer Graphics and Applications

LightPainter: Creating Long-Exposure Imagery from Videos

This article from *IEEE Computer Graphics and Applications*' July/August 2018 issue presents LightPainter, an interactive tool that promotes creative long-exposure photography through an intuitive drawing metaphor and flexible spatiotemporal mapping from videos to composite images. The authors discuss the power of software-defined exposure and the tools needed to create sophisticated long-exposure effects in challenging scenarios.

IEEE Intelligent Systems

Building a Globally Optimized Computational Intelligent Image Processing Algorithm for On-Site Inference of Nitrogen in Plants

Estimating nutrient content in plants is a crucial task in the application of precision farming. This work will be more challenging if it is conducted nondestructively based on plant images captured in the field due to the variation of lighting conditions. The authors of this article from the May/June 2018 issue of *IEEE Intelligent Systems* propose computational intelligence image processing to analyze nitrogen status in wheat plants. They also focus on building a genetic-algorithm-based global optimization to fine-tune the color normalization and nitrogen estimation results. This algorithm is able to enhance the nitrogen estimation results compared to other non-global optimization methods.

IEEE Internet Computing

OmniShare: Encrypted Cloud Storage for the Multi-Device Era

Two attractive features of cloud storage services are the automatic synchronization of files between multiple devices and the possibility of sharing files with other users. However, many users are concerned about the security and privacy of data stored in the cloud. The authors of this article from the July/August 2018 issue of *IEEE Internet Computing* present OmniShare, the first scheme to combine strong client-side encryption with intuitive key distribution mechanisms to enable access from multiple client devices and sharing between users. OmniShare uses a novel combination of out-of-band channels, as well as the cloud storage service itself, to authenticate new devices.

IEEE Micro

Plasticine: A Reconfigurable Accelerator for Parallel Patterns

In this article from the May/June 2018 issue of *IEEE Micro*, researchers from Stanford University describe Plasticine, a new spatially reconfigurable architecture designed to efficiently execute applications composed of high-level parallel patterns. With an area footprint of 113 mm^2 in a 28-nm process and a 1-GHz clock, Plasticine has a peak floating-point performance of 12.3 single-precision Tflops and a total on-chip memory capacity of 16 MBytes, consuming a maximum power of 49 W. It provides an improvement of up to 76.9× in performance-per-watt over a conventional field-programmable gate array (FPGA) over a wide range of dense and sparse applications.

IEEE MultiMedia

Integrating Vision and Language for First-Impression Personality Analysis

The authors of this article in the April–June 2018 issue of *IEEE MultiMedia* present a novel methodology for analyzing integrated audiovisual signals and language to assess an individual's personality. An evaluation of their proposed multimodal method using a job candidate screening system that predicted five personality traits from a short video demonstrates the method's effectiveness.

IEEE Pervasive Computing

Learning from Our Mistakes: Identifying Opportunities for Technology Intervention against Everyday Cognitive Failure

There is growing opportunity for technologies to augment human memory and other cognitive processes, but systems to date typically either address known cognitive impairments such as autism and Alzheimer's disease or look to enhance one's general capacity for a specific task. In contrast to these approaches, other researchers argue that recognition and quantification of human error is key to the design of future computing systems for augmenting the human mind. By focusing on cognitive errors, scientists can first identify frequent, persistent, or severe failures as targets for such systems and then go on to measure the success of any interventions. Read more in the April–June 2018 issue of *IEEE Pervasive Computing*.

IEEE Security & Privacy

Enslaving the Algorithm: From a "Right to an Explanation" to a "Right to Better Decisions"?

As concerns about unfairness and discrimination in "black box" machine-learning systems rise, a legal "right to an explanation" has emerged as an attractive approach for governance. The authors of this article from the May/June 2018 issue of *IEEE Security & Privacy* posit that meaningful information about algorithmic logics is more technically possible than commonly thought, but this exacerbates a new "transparency fallacy"—an illusion of remedy over a substantive provision.

IEEE Software

Software Engineering for Sustainability: Find the Leverage Points!

Software engineers are responsible for the long-term consequences of the systems they design—including impacts on the wider environmental and societal sustainability. However, the field lacks analytical tools for understanding these potential impacts while designing

a system or for identifying opportunities for using software to bring about broader societal transformations. This article from the July/August 2018 issue of *IEEE Software* explores how the concept of leverage points can be used to make sustainability issues more tangible in system design. The example of software for transportation systems illustrates how leverage points can help software engineers map out and investigate the wider system in which the software resides.

IT Professional

Experiments with Ocular Biometric Datasets: A Practitioner's Guideline

Ocular biometrics is a promising research field owing to factors such as recognition at a distance and suitability for recognition with regular RGB cameras, especially on mobile devices. The authors of this article from the May/June 2018 issue of *IT Professional* provide a review of ocular databases available in the literature and discuss diversities among these databases, design and parameter consideration issues during acquisition of databases, and selection of appropriate databases for experimentation.

Healthcare's Paradigm Shift

This issue of *ComputingEdge* features research and think pieces on the topics that matter to you most, from healthcare and cloud computing to graphics and robotics.

Healthcare is experiencing a paradigm shift with the emergence of smart health. Personal health technologies now allow patients to collect a wide range of health-related data outside the doctor's office. In "Harnessing the Power of Patient-Generated Data" from *IEEE Pervasive Computing*, the authors report on the PervasiveHealth 2017 workshop (Leveraging Patient-Generated Data [PGD] for Collaborative Decision-Making in Healthcare)—highlighting current challenges and opportunities, as well as outlining a future research agenda for PGD. In "Silver Bullet Talks with Marie Moe," *IEEE Security & Privacy*'s Gary McGraw interviews a SINTEF research scientist and associate professor at the Norwegian University of Science and Technology about medical-device security—including her own life-saving pacemaker.

The new era of healthcare is also being impacted by newer technologies like blockchain and cloud computing. In "Blockchain: A Panacea for Healthcare Cloud-Based Data Security and Privacy?" from *IEEE Cloud Computing*, the authors study the potential of using blockchain to protect healthcare data hosted in the cloud. *Computer*'s "Is Chocolate Good for You—or, Is the Cloud Secure?" looks at the importance of assurance cases in exploring whether cloud-based solutions are secure.

Another new technology that is having a great impact on the field of computing is machine learning. In *IEEE Computer Graphics and Applications*' "Application of Machine Learning to Computer Graphics," the author provides a sampling of ongoing efforts to incorporate machine-learning modules into existing products and services, as well as new products and services such as image processing and video stylization. In *IEEE Internet Computing*'s "Is Bigger Better When It Comes to Android Graphics Pattern Unlock?," the authors examine whether increasing the Android pattern unlock's grid size could increase the security and complexity of human-generated patterns.

In *IT Professional*'s "ICT: An Emerging Paradigm for Success in Nigeria," the authors discuss how Internet and communications technologies (ICT) can assist people in gaining guidance and counseling to make successful employment inroads, leading to success in finding a job. They specifically discuss such an approach with respect to the growing ICT field in Nigeria.

Finally, Shane Greenstein discusses the basic economics of self-driving long-haul trucks and illustrates some of the gains and challenges from scaling in *IEEE Micro*'s "Adjusting to Autonomous Trucking."

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

Harnessing the Power of Patient-Generated Data

Eun Kyoung Choe
University of Maryland, College Park

Bongshin Lee
Microsoft Research

Tariq Osman Andersen
University of Copenhagen

Lauren Wilcox
Georgia Institute of Technology

Geraldine Fitzpatrick
TU Wien

Department editor:
Sarah Clinch;
sarah.clinch@manchester.ac.uk

The authors report on the PervasiveHealth 2017 workshop, "Leveraging Patient-Generated Data (PGD) for Collaborative Decision Making in Healthcare." They discuss characteristics of PGD, followed by scenarios demonstrating the data-sharing practice among patients, clinicians, and caregivers. The authors also highlight current challenges and opportunities, and outline a future research agenda to envision ways to harness the power of PGD.

We are living in an exciting time that is experiencing a paradigm shift in health and medical sciences. In recent years, personal health technologies have emerged that allow patients to collect a wide range of health-related data outside the clinic. These patient-generated data (PGD) reflect patients' everyday behaviors including physical activity, mood, diet, sleep, and symptoms. Thus, sharing PGD between patients and clinicians can help them communicate about health-related concerns and identify actionable insights. Despite active research in this area, however, it is still unclear how we should go about leveraging these PGD to integrate them into clinical practice.

On 23 May 2017, a group of 16 researchers working at the intersection of health informatics, human–computer interaction, and visualization gathered to discuss challenges and opportunities in leveraging PGD in healthcare. The workshop organizers had three objectives:

- identify contexts where PGD sharing could be useful for patients and clinicians,
- bridge communities working in related fields, and
- distill a set of lessons learned from creating scenarios and working toward solutions as a starting point for future efforts.

The workshop was part of 11th European Alliance for Innovation (EAI) International Conference on Pervasive Computing Technologies for Healthcare (PervasiveHealth 2017).

When planning the workshop, the organizers quickly agreed that looking for a one-size-fits-all solution would be misguided, because the design space of PGD is too broad and complex. Thus, they asked the participants to describe in their position paper a concrete and specific PGD-

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

sharing context they were working on. At the beginning of the workshop, participants presented their PGD research contexts and key challenges in their work. Based on this shared understanding, they collectively delineated a set of dimensions to better characterize PGD, and designed example scenarios considering stakeholders' contexts and goals. In this article, we report back key themes, scenarios, and an underexplored research agenda in PGD sharing contexts that emerged from the workshop discussion.

CHARACTERISTICS OF PATIENT-GENERATED DATA (PGD)

PGD cover a wide range of health-related data that are created, recorded, and gathered by patients, family members, or other caregivers.[1] PGD overlap with patient-reported outcomes (PROs), especially when the PROs reflect patients' health conditions (as opposed to their satisfaction with a healthcare entity, for instance). As PGD data types are diverse, so too are the situations and contexts in which these data can be shared and the approaches to enabling data sharing. Here we describe four PGD dimensions and discuss how they shape specific PGD sharing contexts.

First, *who initiates the tracking*—whether it is a patient, clinician, caregiver, or healthcare entity—influences all aspects of the sharing process: from motivation for collecting data, to the practice of collecting and utilizing the data.[2] If clinicians can guide the patient's tracking throughout the care practice to help ensure data quality, this could increase clinicians' willingness to trust and rely on the data in the long run.

Second, *tracking purpose* necessarily interacts with the first dimension. Researchers have recently identified various reasons why people track personal data.[3] For example, people track data to reflect on past behaviors, be aware of their behaviors, solve a particular problem, or fulfill their curiosity. Additionally, clinicians have different purposes for wanting their patients to track data. For example, they can use PGD to better understand what behaviors are "typical" for the patient and to augment communication in the consultation,[4] as well as to aid in the diagnosis of the patient's condition[5]; they also see tracking as a way to improve patient engagement and health outcomes.[6] Hospitals might use PGD instrumentally for quality monitoring, which may be less useful or visible for patients.

Third, in research contexts, we do not often address *data storage and ownership* issues because the focus of research is often to show a proof of concept. In clinical practice, however, where PGD get stored (whether in a hospital-owned electronic medical record, commercial platform, or patient's notebook) and who directly owns the data affect the practical feasibility and willingness of clinicians to access and leverage PGD.

Finally, different *capture mechanisms* exist. For example, sleep quality can be measured in a subjective manner through self-reporting as well as in an objective and automated manner using a tracking device (for example, restless sleep in minutes). The complementary data from multiple sources allow people to cross-check the reliability of the data and identify insights, although the process of analyzing and comparing the data might be difficult and confusing.

PGD-SHARING SCENARIOS

Scenarios help us understand and analyze how technology could be used to reshape users' activities before a system is built.[7] During the breakout sessions, workshop participants designed three PGD-sharing scenarios, considering a range of contexts and goals in addition to the PGD characteristics described in the previous section. People might have different health conditions, which can in turn lead to different goals: patients with chronic illnesses have different needs than acute patients and healthy individuals. Patients who already have a diagnosis manage clearly defined problems, whereas people with undiagnosed or pre-diagnosed conditions often deal with ambiguous and exploratory problems. Their technology proficiency, age, education level, income level, and access to care could widely vary as well. In the following three scenarios, we describe how

tracking personal data and sharing the data with clinicians might help them address some of the challenges within each scenario.

Scenario 1: Chronic Patient Care of an Older Adult with Low Motivation and Low Technology Literacy

We created this scenario to illustrate an extremely challenging type of patient—not only for clinicians but also for researchers. Sarah is a 68-year-old woman who was diagnosed with type 2 diabetes 15 years ago. She lives alone in a low-income residence located in a suburban city in the US. Although she has access to healthcare, Sarah has limited insurance coverage and worries about her medical bill. She suffers from mild depression and fears the possibility of developing dementia as it runs in her family. Sarah is not familiar with high-tech devices and does not own a smartphone, though she does have Internet access. Her lifestyle is in general sedentary and she has low motivation to manage her health (see Figure 1).

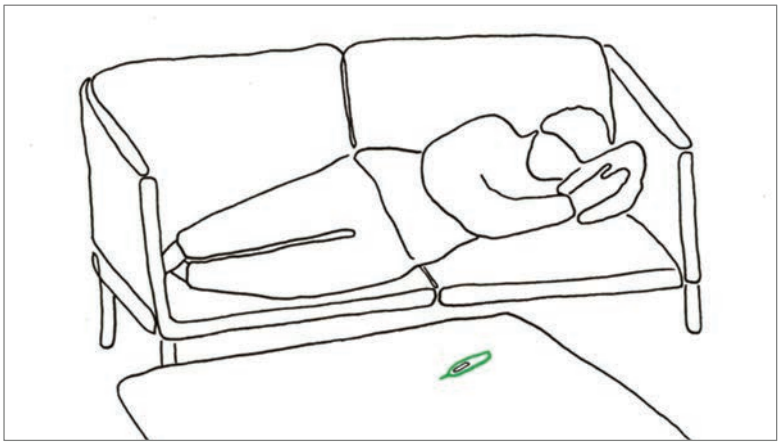

Figure 1. Scenario 1—Sarah, a diabetes patient. Her clinician wants to introduce self-tracking to Sarah as a way to engage her in the care practice.

Sarah dreads experiencing diabetes symptoms such as extreme fatigue and blurry vision as well as more serious complications such as foot problems and kidney failure. Her goal is to avoid these symptoms and complications of diabetes, but it is not clear to her how daily exercise and blood-glucose management connect to these outcomes.

In this scenario, Sarah's goals are not fully aligned with those of her clinician: the clinician wants her to exercise regularly and eat a healthy diet for long-term management, whereas Sarah only wants to avoid the immediate symptoms and complications of her illness. In this case, "low-tech" self-tracking (initiated by the clinician) and data sharing could reduce the gap between the clinician's goals and Sarah's goals. The clinician can introduce self-tracking to Sarah by asking her to start minimal tracking of her exercise and diabetes symptoms; she can enter her duration of exercise (for example, walks) and occurrence of symptoms via a text message or through a tracking website.

They can then use the self-tracking data to discuss the relationship between her day-to-day activities and her diabetes symptoms. In addition, to promote and sustain Sarah's engagement, the clinician can use the data to communicate the importance of tracking in their mutual understanding of Sarah's condition and progress. As Sarah gets used to the tracking, the clinician can also guide her to revise the tracking items. Together they can identify the items that are more important to collect and share, then focus on them; potential data to share include blood-sugar levels, medication taking, diet, mood, stress, and physical activity.

Scenario 2: Casual Data Tracking for Later Use in a Clinical Setting

This scenario is meant to illustrate the potential value of incidental or unpurposeful tracking of "life data" and how they may later become useful for health purposes. Oliver is in his mid-thirties and runs his own hair salon in a large European city. He is trained as a professional hairdresser and has experience from various high-end establishments. Besides his passion for healthy diets and fashion, he finds technology fascinating. Inspired by a close friend, Oliver was one of the first adopters of heart-rate monitors in the 1990s, and ever since he has been exploring a broad range of activity tracking including heart rate, sleep, food, and mood as well as more unconventional measures like gut bacteria and social activity (see Figure 2).

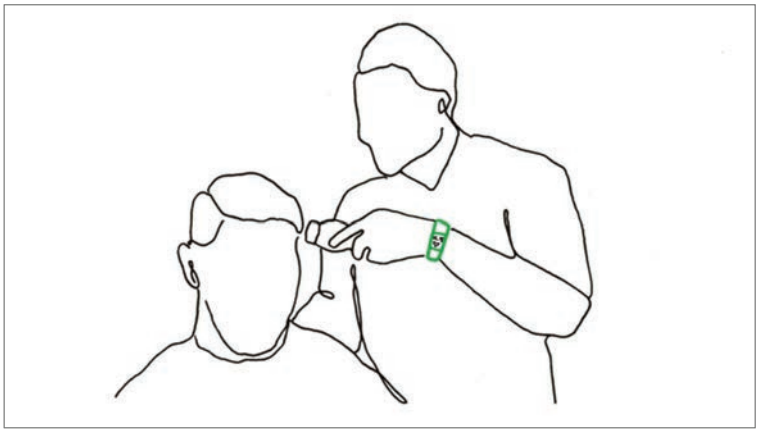

Figure 2. Scenario 2—Oliver, a casual self-tracker. Oliver and his clinician want to utilize Oliver's self-tracking data for a better diagnosis.

His business was going well for a while, but lately sales have started to drop—mainly because of the bad temper and negative attitude Oliver shows toward his customers. Oliver's social life has also changed. His close friends no longer call or write to him and he feels alone most of the time. Left in a life crisis and feeling more and more anxious, he makes an appointment with his doctor. At the meeting, they discuss his social history, how he sometimes loses his temper, and the times he feels anxious and emotionally unstable. Oliver also shows several printouts of some of the PGD he has been tracking over many years.

The doctor cannot reach a diagnosis and suggests that Oliver consult a colleague—a health "informatician" who is skilled in health data analyses and provides services in making sense of large-scale PGD. Upon specific requests, Oliver collects a multitude of different data including self-tracked data from devices and other behavioral data from sources like Facebook, and uploads them to the informatician. At the following consultation, the doctor shows Oliver different examples of data correlations: high blood pressure during periods of low social interaction, high levels of blood glucose associated with bad mood, and increased use of mobile applications correlated with low concentration. Based on the data analysis, the doctor diagnoses Oliver with a mild form of ADHD and begins treatment to improve and restore Oliver's life.

Scenario 3: Parents with a Newborn Baby Capturing and Sharing Health and Developmental Data

This scenario covers a case in which individuals are not capable of tracking their data on their own, and thus caregivers need to capture and share the data on their behalf. Jessica and John are in their early 30s, with a son, Lucas, born just two weeks ago. They both work for IT companies in the US and have good health benefits.

As first-time parents, they have a lot to learn and are keen on keeping a close eye on the baby's health and development. Jessica and John want to figure out basic things including how to feed him and train him for sleep. They are also anxious to know whether the baby is growing according to standard milestones or having any issues. Thus, in addition to logging data (daily feeding, diaper changes, sleep, and so on) in a paper-based diary as the pediatrician suggested, they decide to install a high-end baby monitor on a crib, allowing them to remotely view high-quality video 24/7 (see Figure 3). As it is equipped with an infrared camera, the baby monitor captures Lucas' sleep and wake episodes at night. It also shows room temperature and the baby's day-to-day activities such as turning over, crying, and napping.

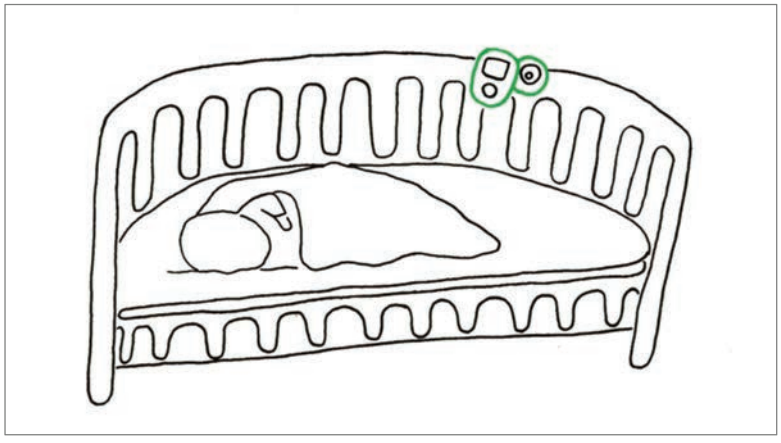

Figure 3. Scenario 3—Jessica and John installed a high-end video camera to capture baby's developmental data, which they hope to share with the pediatrician.

At their well-child visit, during which Jessica and John are encouraged to ask questions about Lucas, they are eager to share these additional data with their pediatrician. One of Jessica and John's main challenges, however, is to share the data and ask questions without overwhelming the pediatrician. To reduce the information overload while increasing the utility of the video data, they want data-capturing tools that help them create a succinct summary from the video. In addition to extracting video clips that capture important developmental activities, they need to be able to annotate the video summary using in situ tagging and ad hoc notes. Once the clinician sees the benefit of the additional video data, she can suggest other markers Jessica and John should look for.

RESEARCH CHALLENGES AND OPPORTUNITIES

As shown in all three scenarios, sharing of PGD between clinicians and patients can affect their personal interactions, workflow, and communication dynamics. While the outlook is positive, we should also attend to research challenges and negative consequences that could emerge in leveraging PGD in medicine.

Supporting Clinicians' Goals through PGD

In all three scenarios, clinicians are key partners in patient–clinician collaboration. As such, it is critical to involve clinicians during the early phase of design discussions, and their needs and goals should be reflected in design requirements and solutions. We developed these scenarios with limited participation from the clinician side, so they do not properly reflect clinicians' situations and perspectives. We assumed that clinicians are willing and motivated facilitators in introducing tracking to patients and leveraging PGD, but they may not be so willing to adopt technology and to discuss PGD with their patients due to workload perceptions and not understanding the possible benefits. Furthermore, they could have low data and visualization literacy,

which is important for effective data-driven communications. To successfully leverage PGD in patient–clinician collaboration, more work is needed to better understand clinicians' goals and desires by working with them.

Facilitating Actionable Insight Gaining from Multiple Data Sources

Collecting personal data from multiple sources has become easier and more prevalent. However, as highlighted in scenarios 2 and 3, these data are scattered across many devices, apps, and platforms,[3] making it difficult to get a holistic and synthesized view of one's health and well-being. Despite the emergence of systems taking integrative approaches to data collection, such as AWARE (www.awareframework.com) and OmniTrack,[8] and visualization, such as Visualized Self[9] and Exist (https://exist.io), further research is needed to find easier ways for lay individuals and clinicians to gain insights from multiple sources of personal data.

Making PGD actionable in the clinic is crucial when improving patient–clinician interaction around self-tracking data. As in the scenario with Oliver, we envision a new role in healthcare, *the informatician*, who can help find actionable insights in heaps of life-data. Informaticians will be capable of using computer power and machine learning to discover and visualize not only correlations but also causations in data. They would play a key role in making incidental, multimodal life-data accessible and useful in the clinical encounter.

Cultivating Sustainable Data Collection

The quality of data plays an important role in data-driven communications. To ensure data quality, we need to better motivate patients to collect their data diligently. First, we can make data tracking meaningful to them. For example, to engage people in tracking and self-reflecting on their fitness data, the Go & Grow system maps individuals' walking data to the amount of water their plant receives.[10] In addition, the possibility of creating a beautiful and unique artifact to express oneself using personal data can draw people to self-tracking. For example, Dear Data (www.dear-data.com) demonstrates that personal data can be visualized in beautiful, creative, and compelling, albeit subjective, ways, encouraging people to diligently collect data.

Ensuring the Clinical Relevance of Collected Data

Clinicians and patients might have very different perceptions about the value of PGD.[1] As it becomes easier to collect large amounts of data that might not be clinically relevant, clinicians are at increasing risk of data overload, which is likely to discourage them to adopt potentially useful PGD. To make PGD useful for clinicians, we need to make it easy and not time-consuming to take action.[11] Given that clinicians can judge the relevance and importance of PGD, there should be a mechanism in place to involve clinicians throughout patients' data-capture process. In this way, clinicians can guide patients to capture markers that are relevant to their disease, hence increasing the value of PGD to clinicians.

CONCLUSION

In health and medical sciences, PGD bring a wealth of research directions and opportunities. In the future, tracking and sharing PGD might change the practice of clinical consultations as we know it: patients and clinicians have a data-driven medical consultation, improving patient engagement and speeding up the diagnosis. We hope to continue discussing ways to leverage PGD for better care with more clinician engagement, and encourage other researchers to contribute to future conferences including PervasiveHealth.

REFERENCES

1. H. Zhu et al., "Sharing Patient-Generated Data in Clinical Practices: An Interview Study," *AMIA Ann. Symp. Proc.*, vol. 2016, 2016, pp. 1303–1312.
2. C.-F. Chung et al., "Boundary Negotiating Artifacts in Personal Informatics: Patient-Provider Collaboration with Patient-Generated Data," *Proc. 19th ACM Conf. Computer-Supported Cooperative Work and Social Computing* (CSCW 16), 2016, pp. 770–786.
3. E.K. Choe et al., "Understanding Quantified-Selfers' Practices in Collecting and Exploring Personal Data," *Proc. SIGCHI Conf. Human Factors in Computing Systems* (CHI 14), 2014, pp. 1143–1152.
4. C. Kelley, B. Lee, and L. Wilcox, "Self-Tracking for Mental Wellness: Understanding Expert Perspectives and Student Experiences," *Proc. 2017 CHI Conf. Human Factors in Computing Systems* (CHI 17), 2017, pp. 629–641.
5. W.J. Korotitsch and R.O. Nelson-Gray, "An Overview of Self-Monitoring Research in Assessment and Treatment," *Psychological Assessment*, vol. 11, no. 4, 1999, pp. 415–425.
6. L. Mamykina et al., "MAHI: Investigation of Social Scaffolding for Reflective Thinking in Diabetes Management," *Proc. SIGCHI Conf. Human Factors in Computing Systems* (CHI 08), 2008, pp. 477–486.
7. M.B. Rosson and J.M. Carroll, *Usability Engineering: Scenario-Based Development of Human-Computer Interaction*, Morgan Kaufmann, 2001.
8. Y.-H. Kim et al., "OmniTrack: A Flexible Self-Tracking Approach Leveraging Semi-Automated Tracking," *Proc. ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies*, vol. 1, no. 3, 2017; doi.org/10.1145/3130930.
9. E.K. Choe et al., "Understanding Self-Reflection: How People Reflect on Personal Data through Visual Data Exploration," *Proc. 11th EAI Int'l Conf. Pervasive Computing Technologies for Healthcare* (PervasiveHealth 17), 2017, pp. 173–182.
10. F. Botros et al., "Go and Grow: Mapping Personal Data to a Living Plant," *Proc. Int'l Working Conf. Advanced Visual Interfaces* (AVI 16), 2016, pp. 112–119.
11. T.O. Andersen et al., "Aligning Concerns in Telecare: Three Concepts to Guide the Design of Patient-Centred E-Health," *J. Computer Supported Cooperative Work*, pending publication.

ABOUT THE AUTHORS

Eun Kyoung Choe is an assistant professor in the College of Information Studies at the University of Maryland, College Park. Contact her at choe@umd.edu.

Bongshin Lee is a senior researcher at Microsoft Research. Contact her at bongshin@microsoft.com.

Tariq Osman Andersen is an assistant professor at the Department of Computer Science at the University of Copenhagen. Contact him at tariq@di.ku.dk.

Lauren Wilcox is an assistant professor in the School of Interactive Computing at Georgia Institute of Technology. Contact her at wilcox@gatech.edu.

Geraldine Fitzpatrick is Professor of Technology Design and Assessment at TU Wien (Vienna University of Technology). Contact her at geraldine.fitzpatrick@tuwien.ac.at.

This article originally appeared in
IEEE Pervasive Computing, *vol. 17, no. 5, 2018.*

Editor: Gary McGraw, gem@cigital.com

Silver Bullet Talks with Marie Moe

Gary McGraw | Cigital

Hear the full podcast at www.computer.org/silverbullet. Show links, notes, and an online discussion can be found at www.cigital.com/silverbullet.

Marie Moe is a security researcher at SINTEF—a multidisciplinary research institute in Trondheim, Norway—and an associate professor at the Norwegian University of Science and Technology (NTNU). She's been a team leader at NorCERT, the Norwegian National CERT, and has taught security classes at Gjovik University College. Her recent work focuses on public safety and security systems that impact human life, including medical device security. Moe received a PhD in information security from NTNU as well as an MSc degree in mathematics.

You're a teacher and a researcher. Which is more fun?

I think they're both fun, and I've been so lucky to get out there and talk to a lot of people and go to conferences. I like talking about my research, doing the research, and teaching future generations of security researchers. So I think everything is equally important.

Which courses do you really enjoy teaching?

I've been teaching a course on incident response and preparedness planning for three years, which will be offered at NTNU. Gjovik University College has merged with NTNU, so some courses are changing. It seems like this will be the last time I give this course in this form, but it's going to continue as a more specialized course for master's students in incident response.

Tell us how the SINTEF lab works and what kind of work you do there.

SINTEF is an independent research institute—it's the biggest one in Scandinavia, with 2,000 employees. My institute is focused on ICT [information and communications technology], and we're doing contracted research for businesses, companies, the government, and the public sector in Norway. There's a small percentage that is funded directly from the Norwegian state, but usually we have to compete for the funding of our projects. Since I started, I've been involved in a lot of application writing to fund research projects.

What work are you most proud of that you're doing at SINTEF?

I'm really proud of the project I helped initiate, which is the Pacemaker Security Project. It started as a hobby project of mine a year ago. I was asked to give a keynote talk at the HACK.lu conference in Luxembourg about being a security researcher who depends on a medical implant—a pacemaker. I've depended on mine for five years as it's correcting my heart rhythm, which was really slow. I passed out suddenly because my heart was taking a break, so the pacemaker saved my life. After getting the implant, I started to look into the device's security.

I had a lot of questions about whether the pacemaker could have any security vulnerabilities or if it could be hacked. Research from 2008 by the University of Michigan and Kevin Fu showed that it can be hacked, but I wanted to know more about the device from my own perspective. I decided that I wanted the people listening to my talk to be aware that there are many insecure medical devices out there and that we need to pay more attention to this problem. I wanted to inspire them

 Published by the IEEE Computer Society

to pick up medical devices and start doing research.

As I was preparing my talk, I thought, "I'm a researcher. Why am I not doing this research myself?" My friend Éireann Leverett and several other researchers agreed to work on this project with me in our spare time, but it wasn't a research project for SINTEF until it started to get a lot of publicity. I was accepted to give a talk at CCC last year, which is the biggest hacker congress in Europe. A Norwegian journalist was in the room and interviewed me, and the article was published in one of the biggest newspapers in Norway. My manager approached me and told me that SINTEF wanted to support the project, so I got some internal funding. I'm writing a grant proposal right now to try to get funding from the Norwegian Research Council.

When you first got your pacemaker, the doctors couldn't answer some of your questions about the technology and how it worked. Do you think you now understand your pacemaker better than they do?

I was admitted to the hospital a couple weeks ago because of a problem with my pacemaker, and at my last checkup, the doctor was talking to me as if I was more of an expert on this than him. I think I might be more knowledgeable about some things on the technological side.

Who do you think should approve modern medical devices—doctors, technologists, or security people? Most doctors don't know about technology, and most technologists don't know much about how the body actually works, so it's tricky.

We need to bring together the different stakeholders, including the device makers, security researchers, regulators, policymakers, standards bodies, healthcare providers, and physicians. Patients should be included in this too.

About Marie Moe

Marie Moe is a research scientist at SINTEF, the largest independent research organization in Scandinavia, and an associate professor at the Norwegian University of Science and Technology (NTNU), where she teaches a class on incident response and contingency planning. She received a PhD in information security from NTNU and was a team leader at the Norwegian Cyber Security Centre NorCERT, where she worked in incident handling of cyberattacks against Norway's critical infrastructure. She's currently researching the security of her own personal critical infrastructure—an implanted pacemaker that is generating every single beat of her heart. Moe lives in Trondheim with her family.

What did you think about the recent claims that pacemakers made by St. Jude Medical have massive security flaws, making them vulnerable to hackers?

I was quite upset when that happened. A lot of people contacted me asking me to comment on it. I was overwhelmed for a week before I decided to write a piece about it in Norwegian for my research team's blog. My first reaction was that I couldn't believe MedSec and Muddy Waters [the firms who wrote the report] didn't contact St. Jude with their findings, and that they didn't contact the FDA [US Food and Drug Administration]. It seems they were financially motivated. I calmed down a bit when I read the report because there were no real details in it that could allow an attacker to easily replicate the attacks. There was a video alleging that the pacemaker had been hacked, and the proof was an error message on the pacemaker's programmer. But Kevin Fu's group showed that you could actually get the same error message by interrogating a pacemaker that wasn't connected to muscle tissue. There was another claim that they could deplete a pacemaker's battery, which is a really bad scenario. [The FDA has since investigated the hacking claims and released a security advisory on 9 January 2017, confirming the vulnerabilities.]

Completely unrelated, St. Jude recently announced a recall for a brand of their implanted cardiac devices—the battery failed so the device ceased to deliver therapy, and two patients actually died because of this failure. What's very problematic in this case is that the Muddy Waters report advised patients to switch off the remote telemetry unit because of the security vulnerability claim, but the medical advice to patients affected by the battery recall is to keep telemetry monitoring on at all times to detect a possibly deadly depletion of the battery.

There have been some very famous software bugs over the years, but it seems like we haven't made as much progress as we need to. Do you think there's been much progress in the field?

I was invited by the FDA to be on a panel last year where they were presenting new approaches to premarket and postmarket risk management of products. Also, medical device vendors have started doing incident reporting and developing vulnerability disclosure policies, so there's actually a way for researchers to contact the companies if they find vulnerabilities. This was lacking previously, but vendors are slowly adopting this and want to be in dialogue with security researchers, which is a good thing. After I

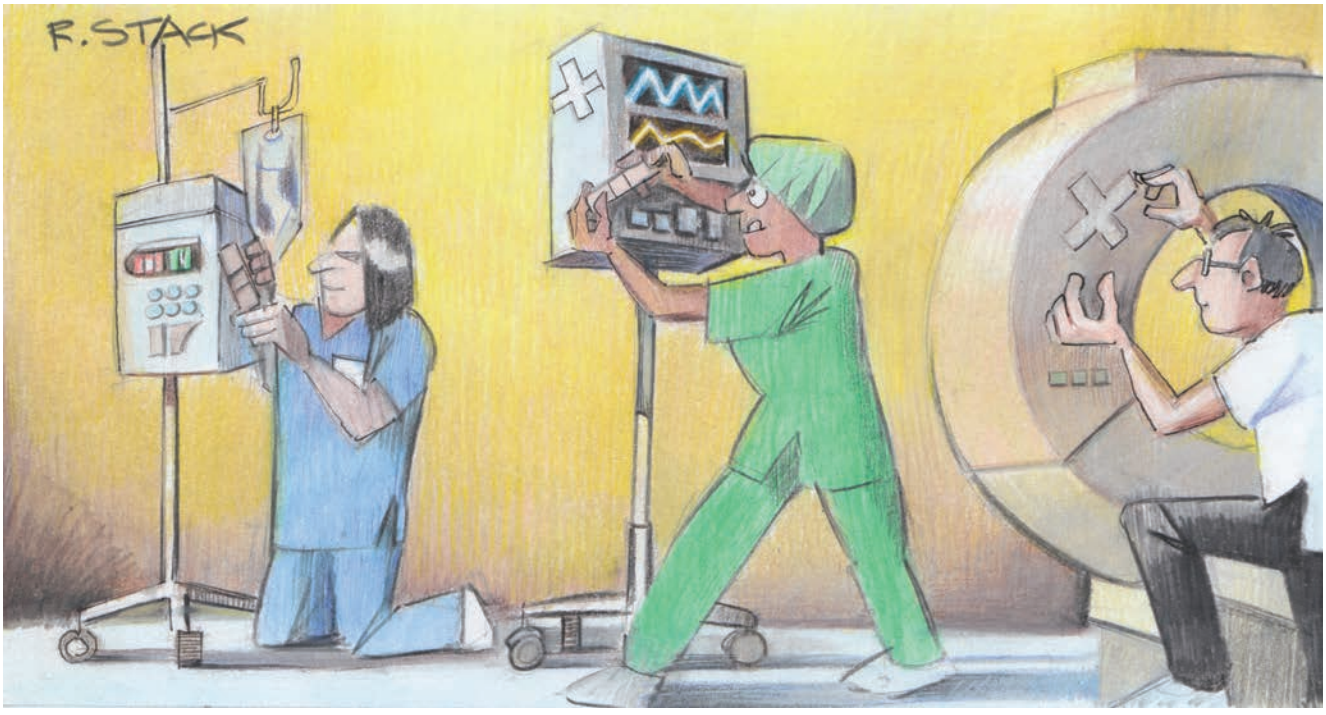

went public with my concerns, I was contacted by several vendors who wanted to talk to me about security. So some progress has been made.

The idea that all you have to do to secure something is hack it, and then fix what you find, doesn't seem right to me. I think you have to design things properly and review the code; it's not just about testing at the end of the lifecycle. What are your thoughts about that?

I completely agree that you need to build security in, and you're the expert on this. Maybe bug bounties aren't the first thing medical device vendors need to focus on. We need more transparency and more testing. I would really like to see more third-party testing of devices. It's the way to go.

You're on the record as wanting medical device code to be open source. Why?

It doesn't have to be open source, but there has to be some transparency that allows end users, patients like me, to trust the products. Of course you can have open source software, but you can also have some alternatives in case that doesn't work. For example, a software bill of materials is a good way to go, because vendors declare all the different pieces of software shipped with the product. For instance, hospitals that don't have a good overview of all their different products can be very vulnerable to ransomware attacks. They need to have their inventory management in place before they can identify what needs patching, but the way it is today, vendors control the devices. They tell the hospitals they need to put holes in the firewalls so they can remotely manage the devices. But the hospitals—not the vendors—are the ones that have to take responsibility for the patients.

Inventory control turns out to be a fundamental aspect of software security. When you think about scalability, you have to do that first.

There's also the issue of putting in measures that can be helpful in an investigation. Today, there aren't any requirements for the forensic capabilities of medical devices, so if something goes wrong, you turn to the vendors to find out what kind of logging they've implemented. You might get crash logs or you might get nothing.

I have a good example of this. Three weeks ago, I was in an airplane on my way to give a talk at a hardware hacking conference, and my pacemaker suddenly failed. I didn't know what was going on at the time, but I could feel my heart beating really hard and I could see that my chest muscle was twitching. I notified the air crew, and when we landed, there was an ambulance waiting that took me to a hospital in Amsterdam. I had to be debugged, in a way. When they hooked me up to the pacemaker programmer, an error message displayed showing that the pacemaker needed a reset. A file was created with the memory dump and logs from the inventory. When my pacemaker was reset, it went back to its factory settings so I needed to have it reprogrammed. Luckily, I had printouts from my previous checkup so I could get my settings back. The next

day, I gave my talk at the conference. The files have been sent to the vendor to investigate this issue.

You survived this because you understand technology, you had a backup, and you knew what your parameters were. But most people don't have that same understanding.

Luckily, I detected that something was wrong and went to the hospital. When my pacemaker failed, it switched into backup mode, which is good because I depend on it to work at all times. The safety mode felt really strange—it had a higher voltage and the electrical current went through my chest-muscle tissue instead of a small closed loop inside my heart—so that was what I detected. But of course, if you don't detect that something's wrong, you might go unaware, which obviously isn't good for your health. I chatted with the pacemaker technician at the hospital, and he said that some devices don't give any error message when things like this happen, so patients don't really know what's going on. It's a problem to not have the logging and the forensic capabilities in place.

How could manufacturers demonstrate to your satisfaction that the medical device they put in your body is secure?

Third-party testing would be a good thing, instead of just trusting that it's secure because the vendor says it is. As a more informed user, I would actually like to see the software myself and do my own third-party testing.

What are your views on the emerging software/device certification marketplace?

I think there's use for software certification, and I think it'll help regulate the device makers and give them incentive to create more secure products.

Let's switch gears. What's it like to be a highly technical woman doing advanced security work in Norway?

There aren't a lot of women in the security field generally, but I actually know a lot of women in Norway who are working in security. At SINTEF my group is actually 50 percent women, which I think is exceptional. I'm used to being one of the few women in the room, and it doesn't really bother me that much. Sometimes it can be to my advantage because people remember me, but at the same time, you need to make a good impression. If I do something wrong, then it's as if all women aren't good at something.

I think that's the really common feeling. In Norway, a yearly conference held by NorSIS, which is focused on women in security, has women speakers and a lot of women attending, which I think is great because it creates a community, and people get to know each other. It's kind of a networking event. I don't know if that's the answer for everything, but at least for students and people starting out in the field, I think it's a good place to go where you don't feel like you're always standing out in the crowd.

The Silver Bullet Podcast with Gary McGraw is cosponsored by Cigital and this magazine and is syndicated by SearchSecurity. ■

Gary McGraw is Cigital's chief technology officer. He's the author of *Software Security: Building Security In* (Addison-Wesley, 2006) and eight other books. McGraw received a BA in philosophy from the University of Virginia and a dual PhD in computer science and cognitive science from Indiana University. Contact him via garymcgraw.com.

This article originally appeared in IEEE Security & Privacy, *vol. 15, no. 1, 2017.*

Blockchain: A Panacea for Healthcare Cloud-Based Data Security and Privacy?

Christian Esposito
University of Salerno

Alfredo De Santis
University of Salerno

Genny Tortora
University of Salerno

Henry Chang
University of Hong Kong

Kim-Kwang Raymond Choo
University of Texas at San Antonio

Editor:
Kim-Kwang Raymond Choo
raymond.choo@fulbrightmail.org

One particular trend observed in healthcare is the progressive shift of data and services to the cloud, partly due to convenience (e.g. availability of complete patient medical history in real-time) and savings (e.g. economics of healthcare data management). There are, however, limitations to using conventional cryptographic primitives and access control models to address security and privacy concerns in an increasingly cloud-based environment. In this paper, we study the potential to use the Blockchain technology to protect healthcare data hosted within the cloud. We also describe the practical challenges of such a proposition and further research that is required.

Healthcare is a data-intensive domain where a large amount of data is created, disseminated, stored, and accessed daily. For example, data is created when a patient undergoes some tests (e.g. computerized tomography or computerized axial tomography scans), and the data will need to be disseminated to the radiographer and then a physician. The results of the visit will then be stored at the hospital, which may need to be accessed at a later time by a physician in another hospital within the network.

It is clear that technology can play a significant role in enhancing the quality of care for patients (e.g. leveraging data analytics to make informed medical decisions) and potentially reduce costs by more efficiently allocating resources in terms of personnel, equipment, etc. For example, data

 Published by the IEEE Computer Society

captured in paper form is hard to capture in systems (e.g. costly and data entry errors), costly to archive, and being available when needed. These challenges may lead to medical decisions not made with complete information, the need for repeated tests due to missing information or data being stored in a different hospital at a different state or country (at the expenses of increasing costs and inconvenience for the patients), etc. Due to the nature of the industry, ensuring the security, privacy, and integrity of healthcare data is important. This highlights the need for a sound and secure data management system.

HEALTH RECORDS IN ELECTRONIC FORMS AND HEALTH INFORMATION SYSTEMS

Generally, *Electronic Medical Records* (EMRs) contain medical and clinical data related to a given patient and stored by the responsible healthcare provider.[1] This facilitates the retrieval and analysis of healthcare data. To better support the management of EMRs, early generations of *Health Information Systems* (HIS) are designed with the capability to create new EMR instances, store them, and query and retrieve stored EMRs of interest.[2] HIS can be relatively simple solutions, which can be schematically described as a graphical user interface or a web service. These are generally the front-end with a database at the back-end, in a centralized or distributed implementation.

With patient mobility (both internally and externally to a given country) being increasingly the norm in today's society, it became evident that multiple stand-alone EMR solutions must be made interoperable to facilitate sharing of healthcare data among different providers, even across national borders, as needed. For example, in medical tourism hubs such as Singapore, the need for real-time healthcare data sharing between different providers and across nations becomes more pronounced.

To facilitate data sharing or even patient data portability, there is a need for EMRs to formalize their data structure and the design of HIS. *Electronic Health Records* (EHRs), for example, are designed to allow patient medical history to move with the patient or be made available to multiple healthcare providers (e.g. from a rural hospital to a hospital in the capital city of the country, before the patient seeks medical attention at another hospital in a different country).[3] EHRs have a richer data structure than EMRs. There have also been initiatives to develop HIS and infrastructures that are able to scale and support future needs, as evidenced by the various national and international initiatives such as the Fascicolo Sanitario Elettronico (FSE) project in Italy, the epSOS project in Europe, and an ongoing project to standardize sharing of EHRs.[4,5,6]

Recently, the pervasiveness of smart devices (e.g. Android and iOS devices and wearable devices) has also resulted in a paradigm shift within the healthcare industry.[7] Such devices can be user-owned or installed by the healthcare provider to measure the well-being of the users (e.g. patients) and inform/facilitate medical treatment and monitoring of patients. For example, there is a wide range of mobile applications (apps) in health, fitness, weight-loss, and other healthcare related categories. These apps mainly function as a tracking tool, such as registering user exercises/workouts, keeping the count of consumed calories, and other statistics (e.g. number of steps taken), and so on.

There are also devices with embedded sensors for more advanced medical tasks, such as bracelets to measure heartbeat during workouts, or devices for self-testing of glucose. For example, Leu and collaborators proposed a smartphone-based wireless body sensor network to collect user physiological data using body sensors embedded in a smart shirt.[8] The data (e.g. user's vital signs) can be continuously gathered and sent in real-time to a smart device, before being sent to a remote healthcare cloud for further analysis. Another example is Ambient Assisted Living solutions for healthcare designed to realize innovative telehealth and telemedicine services, in order to provide remote personal health monitoring.[9]

These developments have paved the way for *Personal Health Records* (PHR), where patients are more involved in their data collection, monitoring of their health conditions, etc, using their smart phones or wearable devices (e.g. smart shirts and smart socks).[10,11]

There are, however, a number of challenges associated with PHRs. For example, can we rely on the data collected by the patients themselves? Should the relevant healthcare providers certified data collected by the patients, and if so, how can this be done? Who should be legally liable for a misdiagnosis or delayed diagnosis, due to decisions being made on the data sent from the patient's device that is subsequently determined to be flawed or inaccurate (e.g. due to a malfunction sensor)?

Despite such challenges and potentially thorny legal issues, having a HIS based on an ecosystem of solutions that is able to seamlessly exchange data among themselves and provide the abstraction of a single health data storage for any given patient (e.g. physically distributed among multiple concrete software instances at multiple healthcare providers and mobile apps) will benefit all users, ranging from patients to healthcare providers to governments.

Cloud computing is a potential solution, due to the capability to support real-time data sharing regardless of geographical locations, to provide resource elasticity as needed, and to handle big data (e.g. hosting of big data analytical tools) to obtain useful insights from the analysis of big healthcare data for research and policy decision making.[12,13]

In Figure 1, we demonstrate how cloud help facilitate sharing of healthcare data among providers, supporting each provider in managing their data, providing a seamless way of exchanging and potentially certifying data between EHR and PHR, and providing a unified/comprehensive view of (the scattered) healthcare records for each patient. In other words, (federated) cloud computing can be used to interconnect the different healthcare providers and their PHR solutions, used by the providers to deal with any sudden or seasonal changes, and so on.

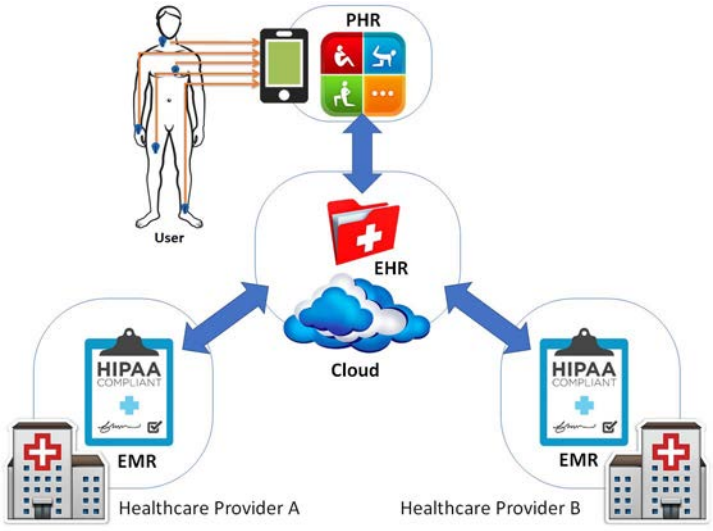

Figure 1. A conceptual cloud-based EMR/EHR/PHR ecosystem.

SECURITY AND PRIVACY

Healthcare data contain personal and sensitive information that may be attractive to cybercriminals. For example, cybercriminals seeking to benefit financially from the theft of such data may sell the data to a third-party provider, who may perform data analysis to identify individuals who may be uninsurable due to their medical history or genetic disorder. Such data would be of interest to certain organizations or industries.

Therefore, ensuring the security of the EMR/EHR/PHR ecosystem and the underlying systems and components that form the ecosystem is crucial, yet challenging due to the interplay and complexity between the systems and components. Moreover, the privacy and integrity of healthcare data must be protected not only from external attackers, but also from unauthorized access attempts from inside the network or ecosystem (e.g. employee of the healthcare provider, or cloud

service provider). The attacks (e.g. leakage or modification of data) can be intentional and unintentional, and organizations may be penalized or held criminally liable for such incidents, for example under the Health Insurance Portability and Accountability Act.

How to secure EMR/EHR/PHR ecosystem and ensure privacy and integrity of the data is an active research area. Approaches include using cryptographic primitives, such as those based on public key infrastructure and public clouds to ensure data confidentiality and privacy.[14] For example, data is encrypted prior to outsourcing to the cloud. However, this limits the searchability of the data, in the sense that healthcare providers have to decrypt the (potentially big) data prior to searching on the decrypted data, resulting in increases in time and costs for the data retrieval and diagnosis (e.g. download, decrypt, and search).[15]

Access control models have also been used to regulate and limit access to the data, based on predefined access policies.[16] Such models can be particularly effective for external attacks, but are generally ineffective against internal attackers as they are likely to be authorized to access the data. There have also been approaches to integrate access control with some cryptographic primitives, such as attribute-based encryption.[17]

BLOCKCHAIN TO THE RESCUE?

There has been recent interest in utilizing blockchain (made popular by the successful Bitcoin) in the provision of secure healthcare data management.[18,19,20] Broadly speaking, blockchain is a technology able to build an open and distributed online database, which consists a list of data structures (also known as blocks) that are linked with each other (i.e. a block points to the following one, hence the name blockchain). These blocks are distributed among multiple nodes of an infrastructure, and are not centrally stored. Each block contains a timestamp of its production, the hash of the previous block and the transaction data, and in our context, a patient's healthcare data and the healthcare provider information.

Figure 2 describes our conceptual blockchain-based EMR/EHR/PHR ecosystem. Specifically, when new healthcare data for a particular patient is created (e.g. from a consultation, and medical operation such as a surgery), a new block is instantiated and distributed to all peers in the patient network. After a majority of the peers have approved the new block, the system will insert it in the chain. This allows us to achieve a global view of the patient's medical history in an efficient, verifiable, and permanent way. If the agreement is not reached, then a fork in the chain is created and the block is defined as an orphan and does not belong to the main chain. Once the block has been inserted into the chain, the data in any given block cannot be modified without modifying all subsequent blocks. In other words, modification can be easily detected. As block content is publicly accessible, healthcare data needs to be protected prior to the data being in the block (e.g. obfuscated and perhaps, encrypted).

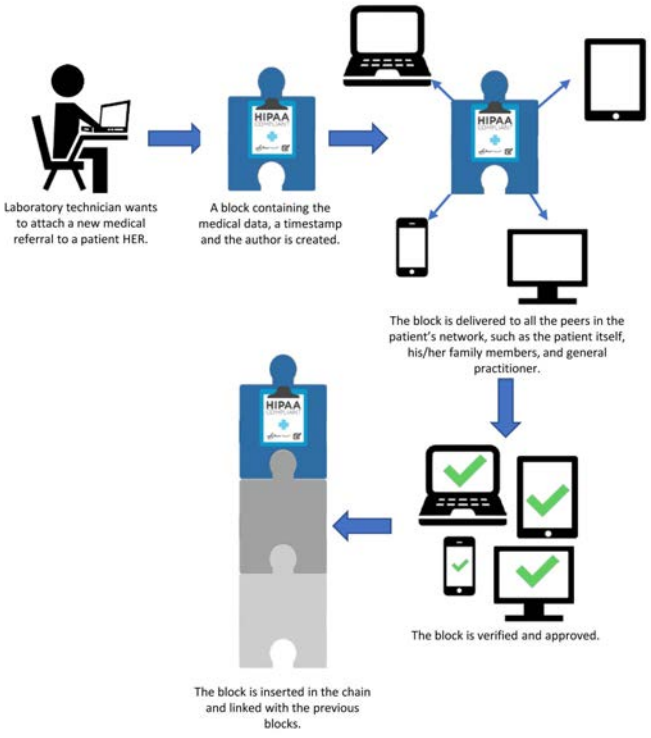

Figure 2. A conceptual blockchain-based EMR/EHR/PHR ecosystem.

Conceptually, blockchain is secure by design that provides the capability to achieve decentralized consensus and consistency, and resilience to intentional and/or unintentional attacks. Key benefits of deploying a blockchain in our approach are as follows:

1. Agreement can be reached without the involvement of a trusted mediator; thus, avoiding a performance bottleneck and a single point of failure;
2. Patients have control over their data;
3. Medical history as a blockchain data is complete, consistent, timely, accurate, and easily distributed; and
4. Changes to the blockchain are visible to all members of the patient network, and all data insertions are immutable. Also, any unauthorized modifications can be trivially detected.

As with any security solutions, there are limitations associated with a blockchain-based approach that need to be carefully studied. For example, blockchain technology can be somewhat disruptive and requires a radical rethink and significant investment in the entire ecosystem (e.g. replacement of existing systems and redesigning of business processes). In other words, before taking the plunge, healthcare providers particularly publicly funded providers will need to undertake a cost benefit analysis to understand the return on investment and any potential implications (e.g. legal and financial). For example, the same record can reside in multiple nodes of the network, located in different countries with different privacy and data protection requirements (e.g. EU and US).

CHALLENGES

While data integrity and distributed storage/access of blockchain offer opportunities for healthcare data management, these same features also pose challenges that need further study.[21]

The strong data integrity feature of blockchain results in immutability that any data, once stored in blockchain, cannot be altered or deleted. However, if the record is healthcare data, then such

personal data would come under the protection of privacy laws, many of them would not allow personal data to be kept perpetually—Article 17 of the soon-enforceable General Data Protection Regulation in the EU has strengthened the rights of individuals to request personal data to be erased. One of the principles of the Organization for Economic Cooperation and Development privacy guideline, on which many data protection laws are based, provides the right-to-erasure to individuals. Given the sensitivity of healthcare data, anyone planning to use blockchain to store them cannot ignore this legal obligation to erase personal data if warranted.

Another practical issue is on how fit it is for blockchain to store healthcare data. Blockchain was originally designed to record transaction data, which is relatively small in size and linear. In other words, one only concerns itself about whether the current transaction can be traced backwards to the original "deal". Healthcare data, such as imaging and treatment plans, however, can be large and relational that requires searching. How well blockchain storage can cope with both requirements is currently unclear.

In order to deal with these challenges, many have suggested the notion of off-chain storage of data, where data is kept outside of blockchain in a conventional or a distributed database, but the hashes of the data are stored in the blockchain. This is said to be the best of both worlds, as healthcare data is stored off-chain and may be secured, corrected, and erased as appropriate. At the same time, immutable hashes of the healthcare data are stored on-chain for checking the authenticity and accuracy of the off-chain medical records.

This idea, however, is not without potential challenges. With the tightening of data protection laws around the world and the attempts by privacy commissioners to regard metadata of personal data as personal data, it may not be very long that hashes of personal data are considered as personal data; then the whole debate of whether blockchain is fit to store personal data may start all over again.

REFERENCES

1. M. Steward, "Electronic Medical Records," *Journal of Legal Medicine*, vol. 26, no. 4, 2005, pp. 491–506.
2. R. Hauxe, "Health Information Systems—Past, Present, Future," *Int'l Journal of Medical Informatics*, vol. 75, no. 3–4, 2006, pp. 268–281.
3. K. Häyrinena et al., "Definition, Structure, Content, Use and Impacts of Electronic Health Records: A Review of the Research Literature," *Int'l Journal of Medical Informatics*, vol. 77, no. 5, 2008, pp. 291–304.
4. M. Ciampi et al., "A Federated Interoperability Architecture for Health Information Systems," *Int'l Journal of Internet Protocol Technology*, vol. 7, no. 4, 2013, pp. 189–202.
5. M. Moharra et al., "Implementation of a Cross-Border Health Service: Physician and Pharmacists' Opinions from the epSOS Project," *Family Practice*, vol. 32, no. 5, 2015, pp. 564–567.
6. S.H. Han et al., "Implementation of Medical Information Exchange System Based on EHR Standard," *Healthcare Informatics Research*, vol. 16, no. 4, 2010, pp. 281–289.
7. D. He et al., "A Provably-Secure Cross-Domain Handshake Scheme with Symptoms-Matching for Mobile Healthcare Social Network," *IEEE Transactions on Dependable and Secure Computing*, vol. PP, no. 99, 2016; doi.org/DOI: 10.1109/TDSC.2016.2596286.
8. F.Y. Leu et al., "A Smartphone-Based Wearable Sensors for Monitoring Real-Time Physiological Data," *Computers and Electrical Engineering*, 2017.
9. M. Memon et al., "Ambient Assisted Living Healthcare Frameworks, Platforms, Standards, and Quality Attributes," *Sensors*, vol. 14, no. 3, 2014, pp. 4312–4341.
10. P.C. Tang et al., "Personal Health Records: Definitions, Benefits, and Strategies for Overcoming Barriers to Adoption," *Journal of the American Medical Informatics Assoc.*, vol. 13, no. 2, 2006, pp. 121–126.
11. S. Marceglia et al., "A Standards-Based Architecture Proposal for Integrating Patient mHealth Apps to Electronic Health Record Systems," *Applied Clinical Informatics*, vol. 6, no. 3, 2015, pp. 488–505.

This article originally appeared in
IEEE Cloud Computing, *vol. 5, no. 1, 2018.*

12. A. Mu-Hsing Kuo, "Opportunities and Challenges of Cloud Computing to Improve Health Care Services," *Journal of Medical Internet Research*, vol. 13, no. 3, 2011.
13. V. Casola et al., "Healthcare-Related Data in the Cloud: Challenges and Opportunities," *IEEE Cloud Computing*, vol. 3, no. 6, 2016, pp. 10–14.
14. S. Nepal et al., "Trustworthy Processing of Healthcare Big Data in Hybrid Clouds," *IEEE Cloud Computing*, vol. 2, no. 2, 2015, pp. 78–84.
15. G.S. Poh et al., "Searchable Symmetric Encryption: Designs and Challenges," *ACM Computing Surveys*, vol. 50, no. 3, 2017.
16. Q. Alam et al., "A Cross Tenant Access Control (CTAC) Model for Cloud Computing: Formal Specification and Verification," *IEEE Transactions on Information Forensics and Security*, vol. 12, no. 6, 2017, pp. 1259–1268.
17. M. Li et al., "Scalable and Secure Sharing of Personal Health Records in Cloud Computing Using Attribute-Based Encryption," *IEEE Transactions on Parallel and Distributed Systems*, vol. 8, no. 3, 2016, pp. 2084–2123.
18. F. Tschorsch and B. Scheuermann, "Bitcoin and Beyond: A Technical Survey on Decentralized Digital Currencies," *IEEE Communications Surveys & Tutorials*, vol. 18, no. 3, 2016, pp. 2084–2123.
19. A. Azaria et al., "MedRec: Using Blockchain for Medical Data Access and Permission Management," *Proceedings of the 2nd Int'l Conference on Open and Big Data* (OBD 16), 2016, pp. 25–30.
20. J. Zhang, N. Xue, and X. Huang, "A Secure System for Pervasive Social Network-Based Healthcare," *IEEE Access*, vol. 4, 2016, pp. 9239–9250.
21. J. McKinlay et al., "Blockchain: Background, Challenges and Legal Issues," *DLA Piper Publications*, 2016; doi.org/https://www.dlapiper.com/en/uk/insights/publications/2017/06/blockchain-background-challenges-legal-issues/.

ABOUT THE AUTHORS

Christian Esposito is an adjunct professor at the University of Naples "Federico II," where he received his PhD in computer engineering and automation. He is also a research fellow at the University of Salerno, Italy. Esposito's research interests include reliable and secure communications, middleware, distributed systems, positioning systems, multiobjective optimization, and game theory. Contact him at christian.esposito@dia.unisa.it.

Alfredo De Santis is a professor of computer science and the department director at the University of Salerno, where he received a degree in computer science. His research interests include data security, cryptography, digital forensics, communication networks, data compression, information theory, and algorithms. Contact him at ads@unisa.it.

Genny Tortora is a full professor of computer science and was dean of the faculty of Mathematical, Natural, and Physical Sciences from 2000 to 2008 at the University of Salerno, where she received a degree in computer science. Her research interests include software-development environments, visual languages, geographical information systems, biometry, and virtual reality. Contact her at tortora@unisa.it.

Henry Chang is an adjunct associate professor at the Department of Law in the University of Hong Kong. His research interests include technological impact on privacy. Chang is a fellow of the British Computer Society and a member of the Hong Kong/Guangdong ICT Expert Committee on Cloud. Contact him at hcychang@hku.hk.

Kim-Kwang Raymond Choo is the holder of the cloud technology endowed professorship in the Department of Information Systems and Cyber Security at the University of Texas at San Antonio. His research interests include cyber and information security and digital forensics. He is a senior member of IEEE, a fellow of the Australian Computer Society, and has a PhD in information security from Queensland University of Technology, Australia. Contact him at raymond.choo@fulbrightmail.org.

IEEE computer society

PURPOSE: The IEEE Computer Society is the world's largest association of computing professionals and is the leading provider of technical information in the field.
MEMBERSHIP: Members receive the monthly magazine *Computer*, discounts, and opportunities to serve (all activities are led by volunteer members). Membership is open to all IEEE members, affiliate society members, and others interested in the computer field.
COMPUTER SOCIETY WEBSITE: www.computer.org
OMBUDSMAN: Direct unresolved complaints to ombudsman@computer.org.
CHAPTERS: Regular and student chapters worldwide provide the opportunity to interact with colleagues, hear technical experts, and serve the local professional community.
AVAILABLE INFORMATION: To check membership status, report an address change, or obtain more information on any of the following, email Customer Service at help@computer.org or call +1 714 821 8380 (international) or our toll-free number, +1 800 272 6657 (US):

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

PUBLICATIONS AND ACTIVITIES
***Computer*:** The flagship publication of the IEEE Computer Society, *Computer*, publishes peer-reviewed technical content that covers all aspects of computer science, computer engineering, technology, and applications.
Periodicals: The society publishes 13 magazines, 19 transactions, and one letters. Refer to membership application or request information as noted above.
Conference Proceedings & Books: Conference Publishing Services publishes more than 275 titles every year.
Standards Working Groups: More than 150 groups produce IEEE standards used throughout the world.
Technical Committees: TCs provide professional interaction in more than 30 technical areas and directly influence computer engineering conferences and publications.
Conferences/Education: The society holds about 200 conferences each year and sponsors many educational activities, including computing science accreditation.
Certifications: The society offers two software developer credentials. For more information, visit www.computer.org/certification.

EXECUTIVE COMMITTEE

President: Hironori Kasahara
President-Elect: Cecilia Metra; **Past President:** Jean-Luc Gaudiot; **First VP, Publication:** Gregory T. Byrd; **Second VP, Secretary:** Dennis J. Frailey; **VP, Member & Geographic Activities: Forrest Shull**; **VP, Professional & Educational Activities: Andy Chen**; **VP, Standards Activities:** Jon Rosdahl; **VP, Technical & Conference Activities:** Hausi Muller; **2018-2019 IEEE Division V Director: John Walz**; **2017-2018 IEEE Division VIII Director:** Dejan Milojicic; **2018 IEEE Division VIII Director-Elect:** Elizabeth L. Burd

BOARD OF GOVERNORS

Term Expiring 2018: Ann DeMarle, Sven Dietrich, Fred Douglis, Vladimir Getov, Bruce M. McMillin, Kunio Uchiyama, Stefano Zanero
Term Expiring 2019: Saurabh Bagchi, Leila DeFloriani, David S. Ebert, Jill I. Gostin, William Gropp, Sumi Helal, Avi Mendelson
Term Expiring 2020: Andy Chen, John D. Johnson, Sy-Yen Kuo, David Lomet, Dimitrios Serpanos, Forrest Shull, Hayato Yamana

EXECUTIVE STAFF

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

COMPUTER SOCIETY OFFICES

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

MEMBERSHIP & PUBLICATION ORDERS
Phone: +1 800 272 6657 • **Fax:** +1 714 821 4641 • Email: help@computer.org
Asia/Pacific: Watanabe Building, 1-4-2 Minami-Aoyama, Minato-ku, Tokyo 107-0062, Japan
Phone: +81 3 3408 3118 • **Fax:** +81 3 3408 3553
Email: tokyo.ofc@computer.org

IEEE BOARD OF DIRECTORS

President & CEO: James Jefferies
President-Elect: Jose M.F. Moura
Past President: Karen Bartleson
Secretary: William P. Walsh
Treasurer: Joseph V. Lillie
Director & President, IEEE-USA: Sandra "Candy" Robinson
Director & President, Standards Association: Forrest D. Wright
Director & VP, Educational Activities: Witold M. Kinsner
Director & VP, Membership and Geographic Activities: Martin Bastiaans
Director & VP, Publication Services and Products: Samir M. El-Ghazaly
Director & VP, Technical Activities: Susan "Kathy" Land
Director & Delegate Division V: John W. Walz
Director & Delegate Division VIII: Dejan Milojičić

IEEE

revised 03 July 2018

Is Chocolate Good for You—or, Is the Cloud Secure?

Kate Netkachova and Robin Bloomfield, City, University of London; and Adelard LLP

In assessing the validity of security claims, assurance cases are an effective approach to help refine the claims, collect detailed evidence, narrow options, and structure a convincing and valid argument to justify the resulting decision.

It's a lovely morning in a university classroom. Early spring sunlight streams through the windows, making the atmosphere bright and cheerful. The smell of chocolate is in the air as the students—security professionals—enthusiastically work on their task. Today is the first day of the Assurance Cases module, and the students have just been introduced to the claims, arguments, and evidence (CAE) framework,[1,2] which is used to develop complex justifications about engineering systems. However, the systems will come later. For now, the students are applying the approach to create a convincing argument about the chocolate that they're eating.

The task is performed in groups, with two different teams working on contradictory claims: "chocolate is good for you" and "chocolate is bad for you." Both teams are positive they can demonstrate their claim's truth. The students actively collaborate to share ideas, split the claim into subclaims, search the Internet for relevant evidence and counterevidence, and develop a strong argument supporting the claim. Along the way, they become more fluent in CAE and gain some initial practical experience in creating structured argumentation.

Finally, it's time to present the results. Both teams come up with a logical structured case supported by reliable evidence, such as research studies and scientific papers, showing that their claim is true.

But how is it possible that two completely opposite claims can both be shown to be true? That's the moment when the students, interested in cloud security, grasp the exercise's purpose.

THE IMPORTANCE OF DETAIL

The chocolate example shows the students, early in the module, the importance of precision in identifying the context, environment, and all related details of a claim before reaching any conclusions:

- *How's chocolate defined?* The amount of cocoa powder, sugar, saturated fat, soy lecithin, and other ingredients can vary substantially, so it's important

 Published by the IEEE Computer Society

EDITOR **JEFFREY VOAS**
NIST; j.voas@ieee.org

to specify its composition prior to developing an argument and collecting supporting evidence.

- *What is the object that "good" or "bad" refers to?* Is it good or bad for one's health, mood, energy level, weight, bank balance, and so on?
- *Who's the claim about?* The health effects of eating chocolate can be very different for someone with diabetes or a nut allergy compared to a perfectly healthy individual.
- *How much chocolate?* The benefits can be quickly outweighed by the risks if more than a moderate amount is consumed, so any assumptions about quantity must be clear.

For engineering systems, asking detailed questions is paramount. The decision to trust a system—for example, to fly in an aircraft or activate a power station—can have physical, societal, environmental, and economic consequences.

Engineering arguments have many characteristics; they're multidisciplinary and science based, with mathematical models and simulations supporting the justification. Overall confidence in any aspect of the system must take into account numerous details. These range from technical considerations (for example, to prove that the software works as intended) to various social aspects (for example, whether we can trust the operators) and is always a judgement made within a particular organizational context to develop a view of the system as whole. The context, environment, system boundaries, and other specific details must be clearly defined before any claim about the system can be demonstrated or any important decision made.

The increasing frequency and severity of cybersecurity incidents and the growing sophistication of attackers bring even more challenges to the decision-making process. Numerous analyses must be performed to address security concerns in addition to the traditional techniques used to assure the system and achieve confidence in engineering decisions.

One way to justify the trustworthiness of complex systems is the *assurance case*, defined as "a documented body of evidence that provides a convincing and valid argument that a system is adequately dependable for a given application in a given environment."[2] Assurance cases support rigorous argumentation but aren't purely deductive—they require inductive reasoning and expert judgement to decide whether they provide sufficient confidence in the decision.

The context, environment, system boundaries, and other specific details must be clearly defined before any claim about a system can be demonstrated or any important decision made.

SECURITY-INFORMED ASSURANCE CASES

The Assurance Cases module we teach at City, University of London focuses on cybersecurity and is delivered within the MSc course Management of Information Security and Risk.

Security analysis is often seen as a distinct activity with its own standards, regulations, culture, and engineering. However, there's a growing realization that security is closely interconnected with other properties and should be integrated into existing analyses rather than performed separately.

Toward this end, we've enhanced the existing assurance-case methodology to analyze and communicate security explicitly.[3,4] In this way, we can consider cybersecurity aspects with other critical system properties in an integrated manner within a security-informed assurance case.

In developing this approach, we found that a significant portion of an assurance case must be changed in the light of security considerations.[4] Incorporating security impacts the design and implementation process as well as verification and validation. Some of the most important security considerations involve system resilience and recovery, in response to malicious events that could change in nature and scope as as the threat environment evolves as well as nonmalicious issues occurring during system operation. Other key aspects include supply-chain integrity and mitigation of the risks of being provided with system components that are compromised or have egregious vulnerabilities and design changes to accommodate user interaction, training, and configuration needs. Addressing software vulnerabilities might lead to additional functional requirements for security controls and new organizational policies.

Security-informed assurance cases show how the technical aspects of security fit within the broader system context, including interdependencies among components and human aspects such as security culture and practices. They provide the means to systematically explore security issues, analyze their impact, and achieve confidence in the security-related decisions made.

Combining security with a wide range of other issues is complicated,

TABLE 1. Security-informed hazard and operability study (hazops) guidewords with examples.

Guideword(s)	Example
No	No message sent
Invalid	Illegal format
Wrong	Wrong data format
Inconsistent	Mismatch between datasets
As well as	Additional message
Other than	Wrong message source, destination
Part of	Element of message missing

and at the heart of this endeavor is a full awareness and understanding of hazards. A detailed hazard analysis must be conducted and integrated into the assurance case to become part of the overall decision-making process.

HAZOPS FOR SECURITY

A hazard and operability study (hazops) systematically identifies potential hazards and deviations from design and operating intention, and is widely used for industrial safety-critical systems.

Security considerations can significantly impact a hazops given the additional risks to system integrity and availability. To bring these areas into focus, we've adapted the traditional hazops to explicitly address security and extended guideword interpretations to facilitate the identification of security-related hazards. Table 1 lists common guidewords with examples.

Prior to performing a security-informed hazops, a simplified architecture diagram is created to capture the most relevant components and system interfaces. Each interface on the diagram is then systematically analyzed by applying the guidewords. In addition to identifying potential accidental hazards and operability problems, as in a traditional hazops, the analysis must also

- identify potential attacks,
- assess credible causes of an attack,
- capture any questions arising about an attack,
- explore the potential consequences of an attack,
- make recommendations to prevent an attack or reduce its consequences, and
- define any needed actions or follow-up activities.

In practice, a security-informed hazops is conducted at multidisciplinary team meetings that bring together the knowledge and expertise of people working on different parts of the system. While performing the analysis, the team explores the entire system. In our capacity as Adelard consultants, we explain the method and guide the process by facilitating discussions and capturing all input. However, the main insight is provided by the system experts; our task is to help them identify security-related hazards and incorporate them into the assurance case.

Adelard's security-informed hazops methodology is based on extensive experience. We've used it to identify and analyze security threats to many complex systems, including large-scale critical infrastructures, and to provide assurance that they're both safe and secure. This is an active research area, and it should be noted there are other related approaches to consider such as the STRIDE threat model[5] and the Cybersecurity Capability Maturity Model.[6,7]

TEACHING COMPLEX CONCEPTS

To deal with the complexity of security-informed assurance cases, we use a pedagogical approach called *spiral learning* (see Figure 1), in which basic concepts are introduced and then repeatedly revisited with more details building upon them.[8] As practitioners, we support hands-on or experiential learning, and research confirms that students learn best by doing.[9–11] We thus introduce various activities and workshops at different stages of the spiral curriculum to ensure that Assurance Cases module participants develop both theoretical knowledge and practical skills.

To facilitate student learning, we use the well-established case-studies approach, which is an effective way to involve students in the experiential learning cycle and increase motivation and interest in the subject.[12–14] We use real-world problems to drive the case studies. The Assurance Cases modules we're conducting this year relate to cloud security. As mature security professionals, many of our students work for organizations that are either using various cloud services now or have plans to do so in the future. Beginning with the deceptively simple question of whether chocolate is good or bad for you, the modules are a useful way to explore whether cloud-based solutions are secure.

HEADS IN THE CLOUD

Cloud computing is a rapidly growing market. According to Gartner,[15] the worldwide public cloud services market is projected to be $246.8 billion in 2017. The highest growth of 36.8 percent will be in infrastructure as a service as adoption becomes mainstream. The software-as-a-service market will see slightly slower growth and is expected to increase by 20.1 percent this year.

Although cloud computing can provide many benefits to organizations, such as greater business agility, scalability, flexibility, and cost optimization, information security is a major concern. With cyberattacks one of the fastest growing economic crimes,[16] companies are increasingly interested in protecting their digital data and minimizing security risks when adopting cloud solutions.

The most common claim students initially make in the Assurance Cases module is that "cloud solution A is secure." However, as with the chocolate example, such a claim is too broad and vague for meaningful analysis.

Security isn't absolute; various details relevant to a particular organization—for example, who uses the service, what devices are used, data types, and usage patterns—must be considered before any conclusion can be made. A more specific claim would be "cloud solution A within organization B provides sufficient level of security for its staff to use for corporate data C on D devices." Starting from this claim, students would need to define "sufficient" security and identify the system boundaries, the specific context and system environment, user and organizational responsibilities and legal obligations, requirements from regulators and other stakeholders, and so on to provide a more detailed analysis.

REACHING A DECISION

Many details will emerge during a security-informed hazops. Students must fully integrate the results of their analysis into the assurance case to ensure their inclusion in the overall decision-making process.

As part of this process, students must clearly distinguish evidence from information. These two concepts are often confused by students, resulting in problems with the justification. Information broadly includes all relevant data, ranging from textual statements to statistics, but only data that specifically supports or undermines a claim constitutes evidence. In an assurance case, it needs to be explicitly clear what each piece of evidence shows and what conclusions it supports.

In addition, students learn not to confuse evidence with claims. For example, one common mistake in assurance cases is to use a service-level agreement (SLA) as evidence for achieving certain requirements. However, SLAs only provide a promise to deliver a given service; they're not actual evidence that the service has been delivered. Therefore, SLAs should only generate claims and require supporting evidence of fulfillment.

With respect to cloud security, unfortunately, students often have trouble obtaining evidence from service providers. Plenty of information is available, but evidence tends to be incomplete or restricted. Nevertheless, students still learn the important skill of determining what evidence is needed.

The cases students produce by the end of the Assurance Cases module are quite detailed. For example, for cloud security cases created by our students in this academic year, the average case contains 37 claims, 25 arguments, and 27 pieces of evidence. Students use CAE blocks[17] with some tool support[18] to structure their cases. On average, this year's student cases contains 10 decomposition, 2 substitution, 2 concretion, and 19 evidence incorporation blocks.

Figure 1. Spiral model for teaching security-informed assurance cases, starting with basic concepts in the center and expanding outward with increasingly complex concepts and additional details.

Returning to this article's title, there's no right answer for such a generic question. It's impossible to demonstrate the truth of "chocolate is good (or bad) for you" or "the cloud is (or isn't) secure." However, there are approaches that can help you reach a decision that satisfies specific requirements within a certain context for a particular application and environment. Built on the CAE framework, assurance cases help our students refine claims, collect detailed evidence, narrow options, and structure a convincing and valid argument to justify the resulting decision.

But assurance cases have a wider impact. In their postmodule evaluations, our students report becoming more cognizant of the claims they and others make and more aware of the need to identify specific evidence, question assumptions, and analyze issues more deeply. This mindset change is precisely what's needed to meet today's increasingly complex security challenges.

ACKNOWLEDGMENTS

This work has been partially supported by the UK EPSRC project "Communicating and Evaluating Cyber Risk and Dependencies (CEDRICS, EP/M002802/1)" and is part of the UK Research Institute in Trustworthy Industrial Control Systems (RiTICS).

This article originally appeared in Computer, *vol. 50, no. 8, 2017.*

REFERENCES

1. *ISO/IEC 15026-2:2011, Systems and Software Engineering—Systems and Software Assurance—Part 2: Assurance Case*, Int'l Org. for Standardization, 2011.
2. R.E. Bloomfield et al., *ASCAD—Adelard Safety Case Development Manual*, Adelard, 1998.
3. K. Netkachova and R.E. Bloomfield, "Security-Informed Safety," *Computer*, vol. 49, no. 6, 2016, pp. 98–102.
4. R. Bloomfield, K. Netkachova, and R. Stroud, "Security-Informed Safety: If It's Not Secure, It's Not Safe," *Software Eng. for Resilient Systems*, A. Gorbenko, A. Romanovsky, and V. Kharchenko, eds., LNCS 8166, Springer, 2013, pp. 17–32.
5. "The STRIDE Threat Model," Microsoft, 2005; msdn.microsoft.com/en-us/library/ee823878(v=cs.20).aspx.
6. *Cyber Security Capability Maturity Model (CMM)*—V1.2, Global Cyber Security Capacity Centre, Univ. of Oxford, 2014; www.sbs.ox.ac.uk/cybersecurity-capacity/system/files/CMM%20Version%201_2_0.pdf.
7. "Cybersecurity Capability Maturity Model (C2M2)," Office of Electricity Delivery and Energy Reliability, US Dept. of Energy; energy.gov/oe/services/cybersecurity/cybersecurity-capability-maturity-model-c2m2-program/cybersecurity.
8. J. Bruner, *The Process of Education*, Harvard Univ. Press, 1960.
9. G. Gibbs, *Learning by Doing: A Guide to Teaching and Learning Methods*, FEU, 1998.
10. H. Wenglinsky, *How Teaching Matters: Bringing the Classroom Back into Discussions of Teacher Quality*, Educational Testing Service, 2000.
11. C. Beard, *The Experiential Learning Toolkit: Blending Practice with Concepts*, Kogan Page Limited, 2010.
12. C. Davis and E. Wilcock, *Teaching Materials Using Case Studies*, UK Centre for Materials Education, 2003.
13. D.A. Kolb, *Experiential Learning: Experience as the Source of Learning and Development*, Prentice Hall, 1984.
14. L.R. Mustoe and A.C. Croft, "Motivating Engineering Students by Using Modern Case Studies," *European J. Eng. Education*, vol. 15, no. 6, 1999, pp. 469–476.
15. "Gartner Says Worldwide Public Cloud Services Market to Grow 18 Percent in 2017," press release, Gartner, 22 Feb. 2017; www.gartner.com/newsroom/id/3616417.
16. B. Sims, "Double-Digit Rise in Crime against UK Corporates as Cyber Becomes Fastest-Growing Form of Economic Infraction," *Risk UK*, 29 Feb. 2016; www.risk-uk.com/double-digit-rise-in-crime-against-uk-corporates-as-cyber-becomes-fastest-growing-form-of-economic-criminality.
17. R. Bloomfield and K. Netkachova, "Building Blocks for Assurance Cases," *Proc. IEEE Int'l Symp. Software Reliability Eng. Workshops* (ISSREW 14), 2014, pp. 186–191.
18. K. Netkachova, O. Netkachov, and R. Bloomfield, "Tool Support for Assurance Case Building Blocks, Providing a Helping Hand with CAE," *Computer Safety, Reliability, and Security*, F. Koornneef and C. van Gulijk, eds., LNCS 9338, Springer, 2015, pp. 62–71.

Application of Machine Learning to Computer Graphics

A Brief Overview

Amit Agrawal
Kleene Closure Consulting

Machine learning (ML) is impacting almost all industries at a rapid pace. The article is a sampling of ongoing ML efforts in the computer graphics industry.

"Artificial intelligence is the new electricity," asserts Andrew Ng, an adjunct Stanford professor who founded Coursera and the Google Brain Deep Learning Project. Professor Ng posits that as electricity transformed every major industry a hundred years ago, machine learning has the potential to transform every industry including computer graphics.

This article captures the current status of this wave as new techniques and technologies are being deployed inside the computer graphics industry. There are two ways that this adoption is happening: (a) incorporation of machine learning modules inside existing products and services, and (b) offering of new products and services that did not exist before.

EXISTING PRODUCTS AND SERVICES

Companies across the landscape are evaluating their pipelines and substituting machine learning components when appropriate in the tool chain. In most cases, these changes are trade secrets that companies keep private. The following is a sampling of some applications of this new capability.

In image editing tools, machine learning algorithms have been used for upscaling, denoising, hole filling, and object replacements. Certain tools like Photolemur (photolemur.com) advertise the use of machine learning, whereas a tool like Adobe Photoshop implements machine learning algorithms behind the scenes.

Among 3D rendering products, AMD's Radeon ProRender and NVIDIA OPTIX 5.0 use machine learning for denoising (see Figure 1 for an example). In addition, NVIDIA is investigating using machine learning for facial animation, anti-aliasing and optimizing light paths for rendering.

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

Figure 1. NVIDIA is using machine learning for denoising rendered images. (Source: NVIDIA; used with permission.)

NOVEL APPLICATIONS

Equally important, machine learning is enabling new products and services that were not possible before.

Image Processing

Magic Pony Technologies (acquired by Twitter, 2016) optimizes transport of game, image, and video assets for web and mobile, by using machine learning techniques for image processing.

Algorithmia (algorithmia.com) is creating a marketplace for ML algorithms, where services are charged based on compute power and royalty per invocation. Algorithmia offers modules for colorizing black and white photos and enhancing the resolution of images, as well as stylization filters.

Last October, Adobe demoed Project Scribbler that colorizes black and white photos automatically as a part of Adobe Sensei. Final details on availability to the users are still forthcoming.

Image and Video Stylization

Products for image and video stylization have existed for some time, but they really took off when the seminal Leon Gatys paper[1] brought the application of deep learning to this problem into the limelight. Some of the applications in this space have had a mobile focus (e.g. Prisma), while others have primarily focused on a web presence (e.g. DeepArt). In either case, the following applications deserve mention: Prisma (prisma-ai.com), Artisto (artisto.my.com), Painnt (moonlighting.io), Style (macdaddy.io/style), and finally DeepArt (deepart.io), which is the productization of Gatys' paper. See Figure 2 for photographs that have been stylized using DeepArt.

Figure 2. DeepArt uses machine learning to stylize photographs to have the style of an artist. (Source: DeepArt.io; used with permission.)

Extensions

There are two extensions worth noting. High resolution stylized images have long been sought after in order to be displayed on large canvas. Since the algorithms are computationally expensive, the workflow usually involves picking a style at a lower resolution and then upscaling while preserving the style. This is usually not straightforward since image up-sampling is unique to the style being used for that image. For an impasto style, the desired detail is the brush marks on the canvas, whereas for watercolor style, the desired detail is the paper texture for the canvas.

Another extension is producing video from still images. Although this poses a variety of problems, the primary challenge is preserving the coherence of the style. In addition, the high cost of computation is prohibitive because the computations have to be repeated 24-60 times per second of video.

According to the author's biased eye, DeepArt does the best job with image quality, whereas Prisma has the best penetration in the marketplace. Although this field faces challenges such as giving users more control over the stylization and storytelling, machine learning has enabled results that were hitherto unachievable.

Real-time Facial Video Stylization

Another area where interesting new offerings have developed is in the field of real-time facial morphing. Looksery (looksery.en.uptodown.com/android, acquired by Snapchat) and MSQRD (msqrd.me, acquired by Facebook) both provide real-time video morphing of the imagery from the front facing camera of a mobile phone. First, the face is tracked from the camera using machine learning. Once a 3D mesh of the face is extracted, the model can be morphed or textured based on the desired effect.

CONCLUSION

In conclusion, this article captures some of the broad areas where machine learning is having an impact on computer graphics. This publication will continue to follow these exciting trends as more of Andrew Ng's predictions are realized.

REFERENCES

1. L. Gatys, A. Ecker, and M. Bethge, "A neural algorithm of artistic style," *Nature Communications*, 2015; https://arxiv.org/abs/1508.06576.

ABOUT THE AUTHOR

Amit Agrawal is the principal at Kleene Closure Consulting. His experience includes leading technical teams creating visual effects for Hollywood blockbusters, running a startup on non-photoreal rendering, using serious games for education, and most recently applying machine learning to the problem of identity. Agrawal received a PhD in Geometric Modeling from the Institute of Robotics and Intelligent Systems at the University of Southern California. He can be reached at amit.agrawal@me.com.

This article originally appeared in
IEEE Computer Graphics and Applications, *vol. 38, no. 4, 2018.*

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

Follow us:

- @ComputerSociety, @ComputingNow
- facebook.com/IEEEComputerSociety, facebook.com/ComputingNow
- IEEE Computer Society, Computing Now
- youtube.com/ieeecomputersociety
- instagram.com/ieee_computer_society

IEEE computer society

Editor: Munindar P. Singh • m.singh@ieee.org

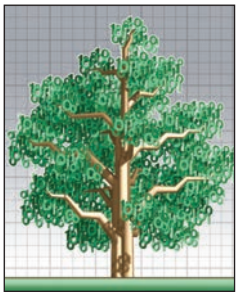

Is Bigger Better When It Comes to Android Graphical Pattern Unlock?

Adam J. Aviv • *United States Naval Academy*

Ravi Kuber • *University of Maryland, Baltimore County*

Devon Budzitowski • *United States Naval Academy*

Android unlock patterns are likely the most prevalent graphical password system to date. However, human-chosen authentication stimuli (such as text passwords and PINs) are easy to guess. Does increasing the grid size from 3 × 3 to 4 × 4 help the situation? Yes and no.

Researchers have proposed a range of graphical password mechanisms,[1] but none have become as prevalent as the Android's graphical pattern unlock. The unlock pattern, or password pattern, is perhaps the most widely used graphical password system to date. Coming preinstalled on all Android phones, it's the choice of a surprising number of Android users.[2]

The password pattern requires users to recall a "pattern" drawn by connecting a set of 3 × 3 contact points arranged in a square grid. A stroke-based pattern must be drawn such that the pattern can be completed without lifting and connect at least 4 contact points, without avoiding any intermediate points. Despite there being 389,112 possible patterns, users select patterns from a much smaller set that are easily guessable, roughly at the same rate as a random 3-digit PIN.[3]

One intuitive and straightforward method for increasing the entropy of user patterns is to increase the number of contact points: Why not allow a 4 × 4, 5 × 5, or even larger grid for users to select patterns? For example, even expanding the grid to 4 × 4 increases the number of possible patterns to 4,350,069,823,024. It's reasonable to expect that expanded grids would also increase the complexity of user chosen patterns.

As part of a series of research papers investigating Android patterns,[4-7] we examined the research question: Does increasing the grid size increase the security of human-generated patterns? We conducted an extensive in-person and online study in support of the research, finding that in some cases, yes — expanded grid sizes can produce stronger patterns, but in many cases, users still choose from a predictable and guessable set. This supports ample anecdotal evidence from password research that humans are poor at selecting strong passwords.[8]

In this article, we review some of the major results of our research and provide insights into future directions.

Data Collection

We conducted two studies to collect Android patterns on both 3 × 3 and 4 × 4 grid spaces. Realism was an important consideration when conducting the research. Although leaked datasets containing real-world, text-based passwords are publicly available, no such resources are available for Android patterns, nor should we expect them to become available. Android patterns are used

 Published by the IEEE Computer Society

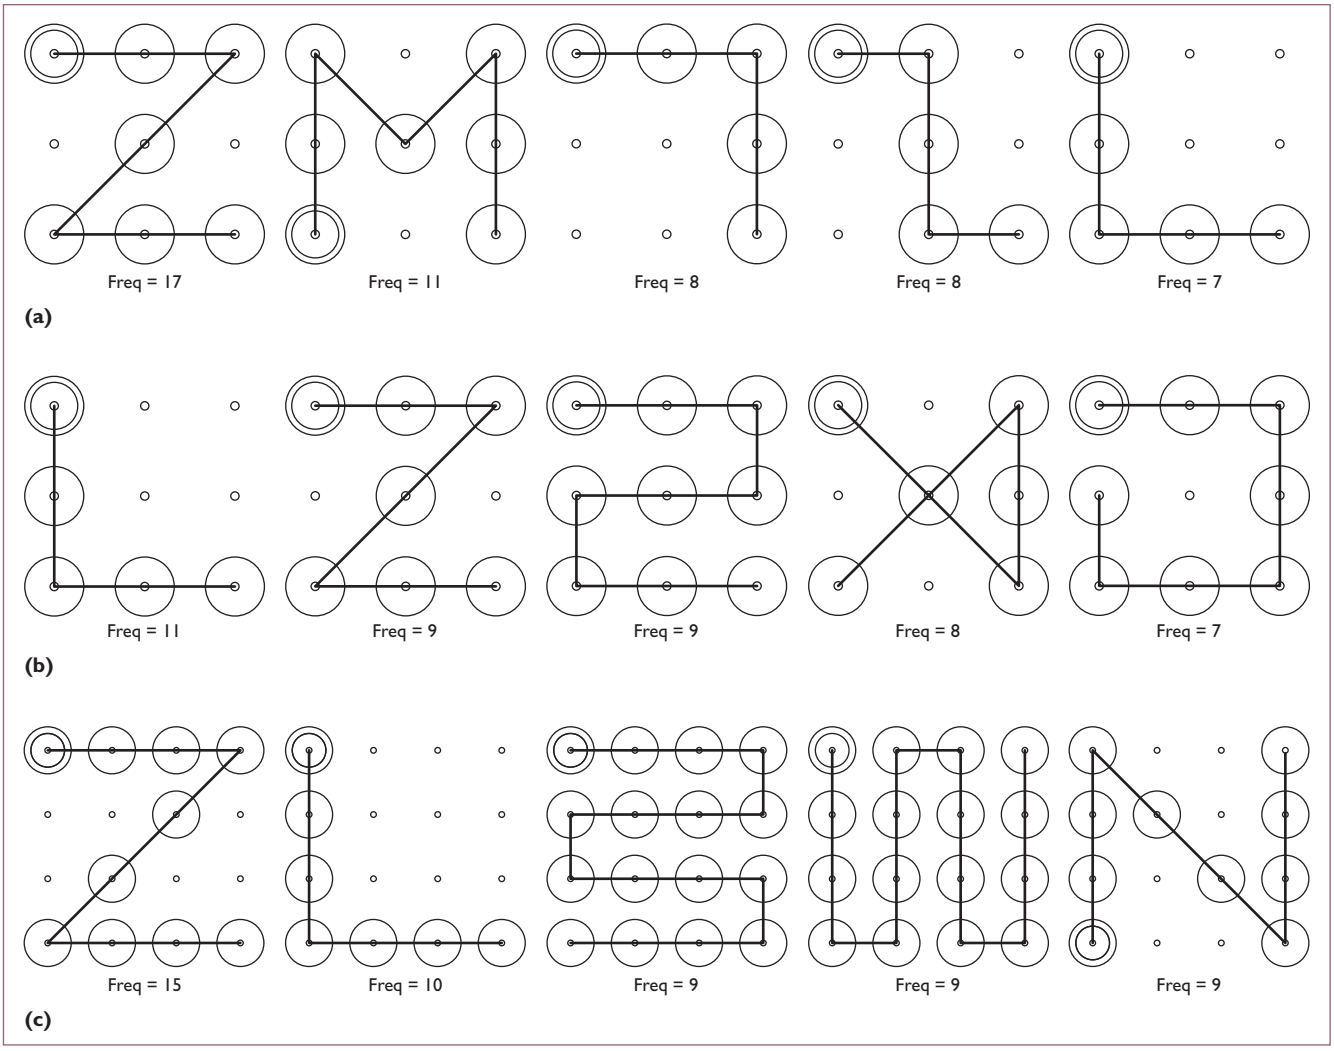

Figure 1. Top 5 most frequently occurring patterns. (a) Self-reported 3 × 3. (b) Pen-and-paper (pen-paper) 3 × 3. (c) Pen-paper 4 × 4.

for local authentication, so there can't be large leaks — and further, the use of 4 × 4 grid patterns is uncommon among Android users. As such, we had to develop new strategies for collecting these patterns.

First, as password patterns are graphical in nature, we developed a pen-and-paper (pen-paper) protocol built upon related work[3] by which participants, literally, draw a set of patterns that they wish others "not to guess" (so-called *defensive patterns*) and then a set of patterns that they believe others choose (so-called *offensive patterns*). If the participant chooses either a pattern that others didn't guess or were able to guess a pattern that others chose, they received a reward.

The pen-paper study was conducted with 80 participants in 10 sets of focus group. These took place over a period of 6 weeks. The participants generated 494 3 × 3 patterns (380 offensive and 114 defensive) and 504 4 × 4 patterns (385 offensive and 119 defensive). Some patterns were rejected from the analysis because they didn't follow pattern-generation rules.

As these drawn patterns might not conform perfectly to real users patterns, we also conducted a large online study of self-reported patterns to compare against the in-person study. Using Amazon Mechanical Turk, we recruited 440 participants who we confidentially asked to report their Android pattern or features of their patterns.

We performed a number of comparisons between the datasets to ensure that self-reported and drawn patterns are similar,[6] finding overall similarity in the datasets in a number of ways, as we'll discuss.

Data Characterization

As a first step in the analysis, we can see in Figure 1 that participants

Table 1. Fraction of repetitions, symmetries, and the embedding of 3 × 3 patterns in 4 × 4 patterns.

Pattern	Size	Repetitions	Symmetries	Embedding
Self-reported 3 × 3	440	203 (46.1%)	336 (76.36%)	N/A
Pen-paper 3 × 3 (all)	491	245 (49.9%)	398 (81.1%)	N/A
Pen-paper 3 × 3 (offensive)	378	187 (48.3%)	309 (79.8%)	N/A
Pen-paper 3 × 3 (defensive)	113	16 (14%)	54 (47%)	N/A
Pen-paper 4 × 4 (all)	501	179 (35.7%)	204 (40.7%)	166 (33.1%)
Pen-paper 4 × 4 (offensive)	382	156 (40.8%)	177 (46.3%)	142 (37.1%)
Pen-paper 4 × 4 (defensive)	119	10 (8.4%)	10 (8.4%)	24 (20.1%)

*Table 2. Statistics of the length measures (mean $[q_1:q_1]$).**

Pattern	Length	Normalized length	Normalized stroke length
Self-reported 3 × 3	6.0 [5:7]	0.7 [0.6:0.8]	2.9 [2:3.5]
Pen-paper 3 × 3 (all)	6.3 [5:7]	0.7 [0.6:0.8]	2.9 [2.2:3.7]
Pen-paper 3 × 3 (offensive)	6.3 [5:8]	0.7 [0.6:0.9]	3.0 [2.2:3.8]
Pen-paper 3 × 3 (defensive)	6.0 [5:7]	0.7 [0.6:0.8]	3.0 [2.4:3.5]
Pen-paper 4 × 4 (all)	9.6 [7:12]	0.6 [0.4:0.8]	3.2 [2.3:3.8]
Pen-paper 4 × 4 (offensive)	9.8 [7:12]	0.6 [0.4:0.8]	3.2 [2.3:3.8]
Pen-paper 4 × 4 (defensive)	8.8 [6:11]	0.6 [0.4:0.7]	3.0 [2.0:3.7]

* The normalized length was calculated by dividing by the total available points, and the normalized stroke length was calculated by mapping the 3 × 3 and 4 × 4 grid on a 1 × 1 Cartesian plane.

tend to choose common, normally shaped patterns. Patterns shaped like letters, such as Z, L, N, or M patterns, are quite common in both 3 × 3 and 4 × 4 data, especially as you consider rotations and transformations, by which a pattern can be rotated or flipped. For example, the Z- and N-shaped patterns are transformations of each other, a so-called *symmetry*.[4]

In fact, large portions of the dataset are symmetric, much more so than one would expect. As Table 1 shows, in some cases 50 percent of the data is symmetric to some other pattern within the dataset. In 4 × 4 patterns, a large fraction of the patterns is simply the embedding of 3 × 3 patterns, where a 3 × 3 pattern is mapped to the 4 × 4 grid space.

Further analysis of the pattern properties, such as length, reveal other similarities between the 3 × 3 and 4 × 4 pattern datasets. Table 2 displays the patterns' length properties, where *length* is the total number of contact points in the pattern. The normalized length is based on the number of contact points, and the normalized stroke length is the length of the lines in the pattern calculated by mapping the grid into a 1 × 1 Cartesian plane.

As you can see, both the distributions for the self-reported and pen-paper 3 × 3 patterns, and the length statistics (particularly the normalized ones) are similar to that of the pen-paper 4 × 4 pattern. This suggests, as apparent in the frequency data, that participants choose patterns of roughly the same properties in the 3 × 3 and 4 × 4 data.

Finally, we can measure the start and end conditions (which contact point a pattern starts with and which point it ends on). As Figure 2 shows, for both 3 × 3 and 4 × 4 patterns, there are strong tendencies for users to begin patterns in the upper-left of the grid, and end in the lower regions, particularly to the right. A random distribution of patterns would have equal likelihood for the start and end points.

All these properties of human-generated patterns — the repetitions, symmetries, and embedding — can all be leveraged by an attacker to inform a guessing strategy to attack

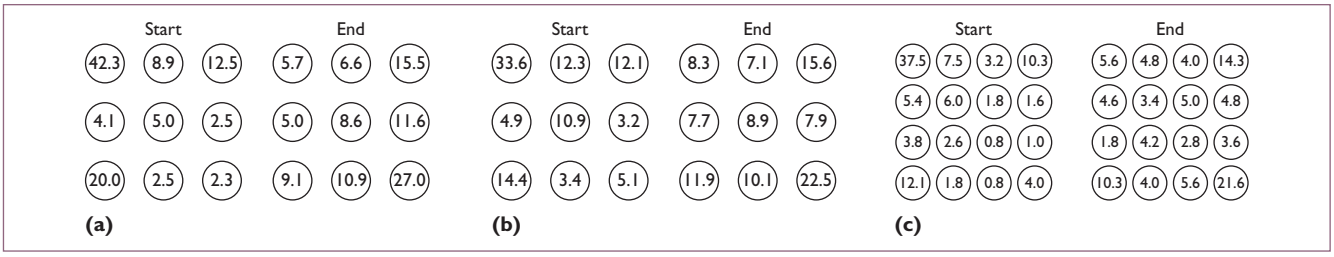

Figure 2. Frequency of pattern start and end points (in percent). (a) Self-reported 3 × 3. (b) Pen-paper 3 × 3. (c) Pen-paper 4 × 4.

the security of graphical patterns, as we performed next to measure the security of patterns.

Measuring Security of Patterns

In the literature of password research, the established method for measuring the security of passwords is based on *guessability*.[9] In its simplest representation, guessability metrics model an attacker's power, asking how many guesses would it take for an attacker, in an offline setting, to *guess* a password. The attacker's knowledge matters; for example, the attacker might have a model of how users select passwords that informs the guesses. The better the model (or the more non-random the passwords being guessed) the easier, or more guessable, the passwords.

We apply a similar notion to the data of graphical password patterns, known as *partial guessing entropy*.[10] Similar to guessability, partial guessing entropy attempts to measure how much randomness (or how hard it would be to guess) a fraction of a set of passwords. This is a common technique employed in prior work[3,11] for analyzing graphical passwords, and allows us to directly compare the relative strength of the different datasets we collected.

Crucial for this metric is developing a method for guessing patterns that accounts for how humans select patterns. Naively, we could simply guess patterns in some arbitrary order, from smallest to largest, based on numbering, and so forth, but this guessing strategy would likely perform poorly because it would take a substantially long time to enumerate the patterns that a human selected, focusing instead on coverage of more complex, less-likely-to-be-selected patterns. Based on a training set, we can instead design a guessing strategy that attempts to leverage the properties described earlier, namely that patterns repeat, are symmetric, and follow regular forms.

Likelihood Estimators

In the development of a guessing strategy, we need to ensure that we don't over-train the data, so we tuned our algorithm using the pen-paper datasets, reserving the self-reported dataset as a *test set* to evaluate the performance of the routine on an independent set.

The guessing algorithm proceeds by assuming that there are some set of sample patterns to train on, and based on that training set, it generates an ordered set of guesses of patterns. We employed a Markov model likelihood estimator to do the training. Using a tri-gram model, we calculated the conditional transition probabilities that two tri-grams were connected in the pattern based on the examples in the training set. The model then calculates a comparable value representing the likelihood of a pattern, which speaks directly to how likely a user would be to choose that pattern.

For example, the Markov model indicates higher likelihood for patterns that start in the upper left and end in the lower right, as indicated by the training data. Further, it can recognize that patterns that include the top row of contact points, moving left to right, are more likely to be followed either by a diagonal set of contact points or a vertical set of contact points, rather than, say, connecting directly to the bottom middle contact point.

There isn't sufficient space to outline the specifics of the Markov model in this article. We refer the reader to our previous work[4] for details of the model. It suffices for this discussion to understand that the model, given a training set and a test pattern, produces a likelihood score. Further, given the model, we can also generate patterns that are likely to occur but might not be directly present in the dataset. These properties will be used in the guessing routine, described next.

Guessing Algorithm

The algorithm uses the training data and the likelihood estimator, trained on that data, to guess a set of patterns. The goal of the guessing algorithm isn't to just guess as many patterns as possible, but as quickly as possible. The order of the guessing directly impacts the partial guessing entropy, which considers what fraction of the dataset is cracked after a certain number of attempts.

As described earlier, to develop the algorithm, we used the pen-paper

*Table 3. Partial guessing entropy comparisons.**

Pattern	$\alpha = 0.1$	$\alpha = 0.2$	$\alpha = 0.5$	Guessed total (%)	Guessed with 20 attempts (%)
Self-reported 3 × 3	**6.62**	**6.95**	**9.49**	**95.9**	**15.0**
Pen-paper 3 × 3 (all)	**6.59**	**6.99**	**8.93**	**97.2**	**16.7**
Pen-paper 3 × 3 (offensive)	6.98	7.69	9.31	95.3	12.5
Pen-paper 3 × 3 (defensive)	9.43	9.79	10.98	90.2	4.0
Pen-paper 4 × 4 (all)	**6.23**	**6.64**	**11.61**	**66.7**	**19.9**
Pen-paper 4 × 4 (offensive)	6.46	7.57	10.40	67.7	16.7
Pen-paper 4 × 4 (defensive)	6.23	6.64	11.61	37.4	3.2
Uellenbeck and colleagues' 3 × 3 (offensive)[3]	7.56	7.74	8.19		
Uellenbeck and colleagues' 3 × 3 (defensive)[3]	8.72	9.10	10.90		
Song and colleagues' 3 × 3 (with meter)[11]	8.96	10.33	12.29		
Song and colleagues' 3 × 3 (without meter)[11]	7.38	9.56	10.83		
Random 3 × 3 pattern ($U_{389,112}$)	18.57	18.57	18.57		
Random 4 × 4 pattern ($U_{4,350,069,823,024}$)	41.98	41.98	41.98		
Random 6-digit PIN ($U_{1,000,000}$)	19.93	19.93	19.93		
Random 5-digit PIN ($U_{100,000}$)	16.60	16.60	16.60		
Random 4-digit PIN ($U_{10,000}$)	13.29	13.29	13.29		
Random 3-digit PIN ($U_{1,000}$)	9.97	9.97	9.97		
Random 2-digit PIN (U_{100})	6.64	6.64	6.64		
Real users' 4-digit PINs[11,12]	5.19	7.04	10.08		

* Here, α refers to the fraction of the dataset being guessed, and results are reported in bits of entropy with comparisons to other related work.

data as the training set, and to test the effectiveness we applied a five-fold cross-validation method on a random selection of 500 patterns from each dataset. The cross validation proceeds by treating four of the five folds as training data, testing on the remaining fold. The results are the average across each result by which each of the folds is treated as the test set.

After much iteration, we developed a guessing routine that best fit our training data. The guessing routine is first informed by the observation that due to the high number of repetitions, the most optimal first step is to guess all unique patterns in the training set, ordered based on repetition frequency, with ties in frequency broken by the likelihood metric. Next, we compute and guess all unique (not previously guessed) symmetries of the training data, again ordered based on the likelihood estimator. Although a large portion of the patterns in the test set are likely to be cracked at this point, we still need to generate more patterns using the Markov model to make 50,000 guesses in total. This process was able to guess 97.2 percent of the human-generated 3 × 3 patterns within 50,000 guesses. Note that there are 389,112 possible patterns.

For 4 × 4 patterns, we can use the same routine, but we can take advantage of additional training information, namely the 3 × 3 patterns. Recall from Table 1 that a large fraction of the 4 × 4 patterns is simply an embedding of 3 × 3 patterns into the larger grid space. So, for 4 × 4 patterns, we included 3 × 3

embedding in the second stage when guessing all the training data's symmetries. After doing so, we are able to crack 66.7 percent of the human-generated 4 × 4 patterns within 50,000 attempts (note that there are 4,350,069,823,024 possible 4 × 4 patterns).

Finally, with the guessing algorithm fixed, we can use all the 3 × 3 pen-paper data as training and attempt to crack the self-reported 3 × 3 patterns — the reserved test set. Table 3 presents the complete results, where we also include the percentage of patterns guessed after 20 attempts, the lockout point on most smartphones.

The results in Table 3 use values in terms of bits of entropy, which we can loosely translate into how much randomness occurs in the dataset (as it relates to how difficult it is to guess that dataset). Note that α refers to the fraction of the dataset guessed, so for example, to guess 50 percent of the self-reported data has 9.49 bits of entropy, which is comparable to guessing a random 3-digit PIN. For 4 × 4 patterns, guessing 50 percent of the patterns has an entropy of 11.61, which is slightly more challenging, but it's less challenging than guessing a 4-digit PIN. Overall, from these results, although there are cases where guessing 4 × 4 patterns are more secure, the common cases aren't significantly more secure than 3 × 3 patterns, suggesting that the expanded grid sizes aren't as beneficial for security as you might imagine.

This work shows that the impact of humans on security systems can be severe. Despite the fact that there are trillions of possible 4 × 4 patterns, the same habits of 3 × 3 patterns persist, greatly weakening the security system. As a larger trend, we as a community need to look for methods that improve security not just through increased choices, but with better directions and instructions. Users don't understand the power of computers to crack passwords, but there are straightforward and simple procedures for improving security — for example, simply picking a pattern that doesn't start in the upper left makes a huge impact. In the same way, for other systems, we need to consider the impact of humans on security systems; humans are often the weakest link.

Acknowledgments

This work was supported in part by the US Office of Naval Research and the US National Security Agency, and we were assisted by Jeanne Luning-Prak, Flynn Wolf, and Rida Bazzi.

References

1. R. Biddle, S. Chiasson, and P.C. Van Oorschot, "Graphical Passwords: Learning from the First Twelve Years," *ACM Computing Surveys*, vol. 44, no. 4, 2012, article no. 19.
2. S. Egelman et al., "Are You Ready to Lock?" *Proc. 2014 ACM SIGSAC Conf. Computer and Comm. Security*, 2014, pp. 750–761.
3. S. Uellenbeck et al., "Quantifying the Security of Graphical Passwords: The Case of Android Unlock Patterns," *Proc. 2013 ACM SIGSAC Conf. Computer & Comm. Security*, 2013, pp. 161–172.
4. A.J. Aviv, D. Budzitowski, and R. Kuber, "Is Bigger Better? Comparing User-Generated Passwords on 3 × 3 vs. 4 × 4 Grid Sizes for Android's Pattern Unlock," *Proc. 31st Ann. Computer Security Applications Conf.*, 2015, pp. 301–310.
5. A.J. Aviv et al., "Smudge Attacks on Smartphone Touch Screens," *Proc. 2010 Workshop on Offensive Technology*, 2010, pp. 1–7.
6. A.J. Aviv, J. Maguire, and J.L. Prak, "Analyzing the Impact of Collection Methods and Demographics for Android's Pattern Unlock," *Proc. Usable Security Workshop*, 2016; www.usna.edu/Users/cs/aviv/papers/aviv-usec16.pdf.
7. A.J. Aviv et al., "Practicality of Accelerometer Side Channels on Smartphones," *Proc. 28th Ann. Computer Security Applications Conf.*, 2012, pp. 41–50.
8. A. Adams and M.A. Sasse, "Users Are Not the Enemy," *Comm. ACM*, vol. 42, no. 12, 1999, pp. 40–46.
9. P.G. Kelley et al., "Guess Again (and Again and Again): Measuring Password Strength by Simulating Password-Cracking Algorithms," *Proc. IEE Symp. Security and Privacy*, 2012, pp. 523–537.
10. J. Bonneau, "The Science of Guessing: Analyzing an Anonymized Corpus of 70 Million Passwords," *Proc. 2012 IEEE Symp. Security and Privacy*, 2012, pp. 538–552.
11. Y. Song et al., "On the Effectiveness of Pattern Lock Strength Meters: Measuring the Strength of Real World Pattern Locks," *Proc. 33rd Ann. ACM Conf. Human Factors in Computing Systems*, 2015.
12. H. Kim and J.H. Huh, "PIN Selection Policies: Are They Really Effective?" *Computers & Security*, vol. 31, no. 4, 2012, pp. 484–496.

Adam J. Aviv is an assistant professor of computer science at the United States Naval Academy in Annapolis, Maryland. His research interests include usability and security, particularly related to mobile devices. Aviv has a PhD in computer and information science from the University of Pennsylvania. Contact him at aviv@usna.edu.

Ravi Kuber is an associate professor of information systems at the University of Maryland, Baltimore County. He's worked extensively in accessibility research and computer-human interaction. Kuber has a PhD in information systems from Queen's University Belfast. Contact him at rkuber@umbc.edu.

Devon Budzitowski is a lieutenant (junior grade) in the US Navy. He's working on his MSE in computer science at the Naval Postgraduate school in Monterey, California. Budzitowski has a BS in computer science from the US Naval Academy. Contact him at debudzit@nps.edu.

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

EDITOR: Gustavo Rossi, Universidad Nacional de La Plata, gustavo@lifia.info.unlp.edu.ar

ICT: An Emerging Paradigm for Success in Nigeria

Bernard Ijesunor Akhigbe, *Obafemi Awolowo University, Nigeria*
Oyinkansola Onyinyechi Akhigbe, *University of Ibadan, Nigeria*
Ishaya Peni Gambo and Babajide Samuel Afolabi, *Obafemi Awolowo University, Nigeria*

A career is an occupation, undertaken for a significant time period in a person's life, that offers opportunities for progress. This implies that—in this context—success is progress. It also suggests that people should choose careers to which they can really give their best. Humans have myriad life choices. At every point in our lives, we must make a lot of decisions. Moreover, innate in every human is self-interest; in the context of choosing a career, we can view such self-interest in relation to those beneficial life opportunities an individual believes can be acquired through career actions based on specific educational pursuits. Career choices, when made correctly, can provide individuals with an image of who they are and what they are capable of.[1]

Furthermore, humans are psychological and cognitive beings. Settling on a specific career and following through with that path can be difficult because there are many choices to make. Humans' mental instincts, which are steeped in emotion, can be so deep that the desire for wider consultations and counseling seems inevitable. People fall prey to confusion when making choices. Noesis is strong in humans—that is, we can reason and perceive things and make comparisons between what we've learned and perceived. However, this innate ability to come up with alternatives, each alternative with its own reward for us to pursue, highlights the complexity for humans in deciding what to pursue that will ultimately make us successful in life. In other words, it can be quite difficult for us to settle easily on a particular life choice or career. A good question is thus how we can seek help in knowing (or being assured of) what to commit our entire lives to and be assured of progress.

When we consider satisfaction with life choices, the foregoing becomes even more complex. As mentioned, we all see or perceive things differently. We might be in the same class and discipline and make grades good enough to show that we will be good at a particular career. However, when the issue of interest and what a person derives satisfaction from comes into the picture, the range of choices about what to pursue can become hazy. This is likely what drives the need for *guidance and counseling* (G&C). This old tradition has certainly thrived over the years. But the fuzzy nature of decision-making makes it paramount that we seek a paradigm that can assist in solving the innate puzzles that often becloud humans when it comes to life choices.

ICT has significant capacity and myriad capabilities to assist in this realm. Human counselors, and even counselees (on their own), can use ICT for guidance. The dicey experience of seeking answers to what an individual should put all his or her energy into throughout life could be settled. For example, it has been argued that for people to make successful employment inroads, they will need not only the know-how but also the psychological make-up to function in a turbulent job market that changes constantly in both

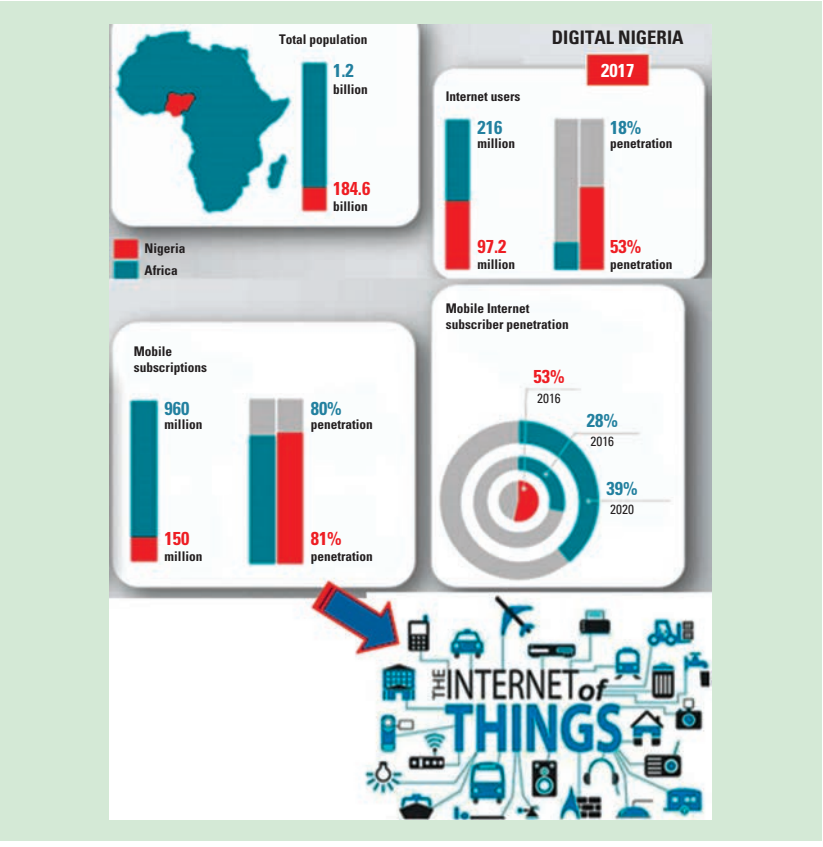

Figure 1. Nigeria's level of ICT savvy as compared with the rest of Africa. (Source: Guitton, 2017[3])

boundaries and requirements. This dynamic might not be easily addressed through traditional G&C.

A paradigmatic approach that is thematic and stochastic as well as unbiased, which might not be possible with the traditional approach, is therefore needed. This approach should also be able to delve into the nitty-gritty of humans' behavioral roots. This position considers the significant potential ICT offers. For instance, ICT is a leading factor in the promotion of creativity and result-oriented processes based on experimentation. This often results in value chains across industry and service delivery. ICT's influence in this regard is felt even now in all areas of human endeavor, including health and social care, the modernization of services in public interest domains, the environment, and so on. For policy formation and the transparency (and accessibility) of governance, ICT is indispensable. However, with respect to providing career and life choice G&C, more work is needed to fully integrate ICT as a tool for pushing the frontier of success in Nigeria.[2]

Background

Nigeria's sovereignty dates back to 1 October 1960. Based on the latest UN estimate, its current population is about 192,692,320 as of September 2017 (bit.ly/2doaI4e). With a federation of 36 states and a federal capital territory, Nigeria holds approximately one-sixth of Africa's population, and is thus the continent's most populous nation. This statistic presents a rationale for considering people's pursuit of career and life choices. Although Nigeria's population should be an advantage in terms of its volume of business transactions, several millions out of the estimated population live below poverty lines. The majority of this population has no clue about what career to pursue so as to earn a living.

Internet usage has increased astronomically since its introduction into the polity. With the proliferation of ICT such as mobile phones, Internet penetration has increased to about 53 percent. This is the highest in Africa, with an estimated mobile subscription rate of 81 percent. This makes Nigeria not just an Internet-savvy country, but one that's savvy to the Internet of Things.[3] The analytic results from comparing Africa's overall Internet savviness with that of Nigeria are shown in Figure 1. We contemplate the factors of Internet (that is, ICT) savviness and population because of the possibility of leapfrogging development phases with respect to ICT and IT. This bodes well for the country, because gains from ICT and IT will not only be realizable in Nigeria, but will even be in compliance with millennium development goals.

In the Nigerian educational sector, there are well-defined and structured manual approaches to career G&C. However, there is no known career G&C framework that uses or even suggests using ICT to assist human counselors. Generally, it has been acknowledged by educationists in Nigeria that decisions about career and life choices can be dicey. Such choices are important because they can determine a person's

prospects for success in life. In Nigeria, the choice has become more complex—as a developing country, there are serious economic problems resulting from unemployment. Current developments in science and technology that can result in new careers, and people's inability to assess themselves to make realistic career choices, are influential factors.

Other factors that have made career choice complex in Nigeria include a lack of information about occupations and job satisfaction. This can inadvertently cause people to make erratic decisions and as such be out of a job.[4] So, ICT's function as a paradigm for providing accurate counsel that will guide those in need about which career to pursue is important. Humans' behavioral twists are ever-changing. This means that we must consider several factors in multiple dimensions,[5,6] thus requiring more sophisticated paradigms to harness the factors that foretell what one will be good at in terms of which career to pursue in life. These are logical motivations to encourage the use of ICT as a paradigm for pushing Nigeria's success frontier from the perspective of career and life choices.

The Potential of ICT

Today, there are ICT techniques that can reduce the partiality that prevents objective consideration of a situation. This could well be necessary when considering sensitive issues such as what people should commit their lives to. As it is, *machine learning* (ML) abstractions can be implemented to reduce bias. ICT remains a suitable tool for encouraging individualization and satisfaction because it can be manipulated and domesticated to guide users in making life choices. ML techniques (MLTs) can learn without being explicitly programmed. These days, even computational statistics are useful for making predictions. That is, computers can be used to deliver mathematical optimization methods, theory, and applications to solve real-life problems.

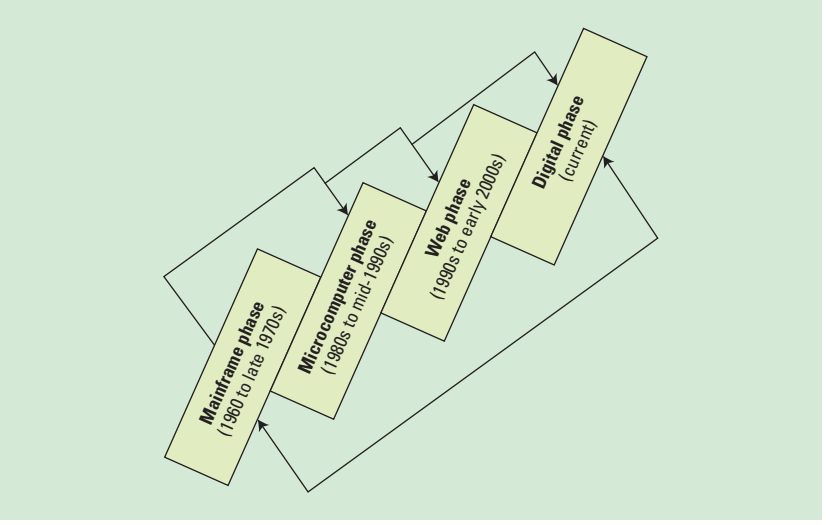

Figure 2. The four phases of ICT development. (Source: Iacob, 2012[8]).

Data analytics using MLTs are useful in devising complex models and algorithms that lend themselves to prediction based on reliable, repeatable decisions and results. This means that ICT can learn from baseline behavioral profiles (data) of entities (such as humans) to find meaningful patterns. In our opinion, this provides endless opportunities, and—in the context we discuss here—can help to guide career path and life choices that cannot be made easily by individuals. Moreover, the iterative aspect of ML is important to consider—that is, as ML models are exposed to new data, they can independently adapt. It has been argued that although ML is a science that is not new, it is currently gaining fresh momentum (see bit.ly/2xCJ4EN, bit.ly/2xBpEVj, and bit.ly/1I5qib7).[7]

Sub-MLTs such as artificial neural networks, regression, and decision trees use a connectionist approach inspired by humans and biological neural networks. Systems developed using such sub-MLTs learn from progressively improved performance to do tasks by considering examples. Supervised MLTs, such as the Apriori algorithm, *k*-means, and so on, consist of two variables: an outcome variable that is dependent and the predicted variable from a given set of predictors. With these variable sets, a function maps inputs to desired outputs, continuing until the model achieves a desired level of accuracy on the training data. For unsupervised MLTs, such as the Markov decision process, there is usually no target variable to predict (or estimate). For reinforcement MLTs, the algorithm is used to train machines to make specific decisions. This makes it possible to expose ICTs to environments in which they train themselves continually based on past experiences. These ICTs try to capture the best possible knowledge to make accurate decisions.

The ever-changing dynamics in human nature and a penchant

for what will bring us value, fame, respect, and some form of influence necessitate the need for career and life choices. This calls for techniques and tools that will assist psychologists, teachers, and amateur and professional G&Cs. ICT as it stands could help realize cutting-edge approaches to mentoring and guiding people through the herculean task of deciding what to spend the rest of their lives on in terms of a career. Historically, there are four identified phases in the development of ICT that we can apply to guidance (see Figure 2). Three trends—increased accessibility, increased interactivity, and more diffused origination (the trend toward more diverse creators and career providers)[8]—are identified within these phases. It will be easy to leverage the practices portrayed, particularly in Nigeria's current phase. In this age of information, Nigeria can leapfrog over the first three phases thanks to ICT's ability to overcome many obstacles (see bit.ly/2xCJ4EN, bit.ly/2xBpEVj, and bit.ly/1I5qib7).[7]

Note also in Figure 2 the rationale to establish the fact that ICT is a paradigm of priority due to its characteristic trend of increased accessibility, interactivity, and more diffused origination. The case of origination provides assurance that there will be more diverse ICT creators and providers driving career and life choice guidance if given the right attention.

Factors that Influence Career and Life Choices

When people are personally satisfied with work and have interest in a particular job, this can intrinsically influence their career choice. To enter a field of endeavor, we must also consider our level of autonomy and creativity. Many other factors, including a lack of competence and confidence to pursue a career, can result from feelings of being inferior or intimidated. ICT offers a wide range of career guidance tools encompassing three main purposes: first, it is a resource; second, it is a medium for communication; and third, it enables the development of materials. By focusing on ICT as a resource for making life and career choices, the possibilities it offers can be used to harness influential factors.

Thanks to the growing use of the Internet and other related ICT, which has arisen due to the proliferation of mobile phones, Nigerians are leaving digital trails—that is, various factors that can help to determine people's behavioral patterns. These patterns represent the intrinsic and extrinsic factors ICT can learn from to help guide a prospective career enquirer. Common among these ICTs are smartphones, via which people carry out several digital activities that generate a lot of patterns. These patterns can be learned by machines and the resultant models implemented in ICT devices that will help provide succinct career guidance. ICT is also important for career guidance and counseling because career work is embedded across a range of public and private contexts.[2] This variety cuts across educational and training provisions to outplacement. The diversity of careers makes it important to consider ICT as an important paradigm for handling such diversity, even when it is reflected in a nomenclature that includes career guidance, counseling, mentoring, coaching, advice, and information dissemination. Furthermore, one-to-one career interviews are difficult to standardize, and their impact is notoriously difficult to quantify.[9,10] Using an ICT paradigm, we can solve all these challenges in Nigeria.[8]

As described (see Figure 1), Nigeria is poised to use ICT in sundry areas of human endeavor. ICT's adoption will no doubt contribute to repositioning the tradition of career G&C both in the educational sector and in real life. This is because the stage is set; a robust and quality ICT infrastructure is available. What's next is to promote strict compliance with policies that encourage a high adoption of ICT, particularly to help both early and late starters seeking a career to determine where to focus to find fulfilment.

Interestingly, this article is one of the few works on using ICT to help people find and follow a career path. Despite the initial bureaucracy that delayed the acceptance and use of ICT to solve problems in Nigeria, we believe that it is still possible to catch up to developed nations. This is because sharing and accessing information is no longer something ICT practitioners need only imagine how to enforce. Due to the proliferation of ICT-related devices such as smartphones, the foregoing has become an everyday reality. Therefore, it is not a matter of whether Nigeria can adopt ICT as an emerging paradigm to push for success in career and life choices. The challenge is rather how best to use ICT like traditional career G&C in a more successful way that will surpass the evidence-and-result-based practice traditional G&C has lived up to. Usually, school G&C programs have two goals: to promote the developmental assets of students and to reinforce positive factors that build resilience.[11,12]

This article originally appeared in
IT Professional, *vol. 19, no. 6, 2017.*

These underlying ingredients must be determined in individuals to be able to encourage them into a particular career or life path. We believe, based on what we've presented about the usefulness of adopting ICT similar to traditional G&C, the foregoing two goals will be achieved. IT

References

1. A. Kim, "The Curious Case of Self-Interest: Inconsistent Effects and Ambivalence toward a Widely Accepted Construct," *J. Theory of Social Behavior*, vol. 44, no. 1, 2014, pp. 99–122.
2. N. Tandon et al., *A Bright Future in ICT Opportunities for a New Generation of Women*, ITU report, Feb. 2012; bit.ly/2kLPzno.
3. B. Guitton, *A White Paper on Nigerian Mobile Trends 2017*, Apr. 2017; bit.ly/2sbTw7N.
4. M. Omoegun and B. Buraimoh, "Career Guidance for Nigerian Students: Why Career Choice Is Becoming More Difficult," CDNet, 2017; cdnetng.org/?q=node/4006.
5. B.I. Akhigbe, B.S. Afolabi, and E.R. Adagunodo, "Modeling User-Centered Attributes: The Web Search Engine as a Case," *Knowledge Organization*, vol. 42, no. 1, 2015, pp. 25–39.
6. B.I. Akhigbe, S.O. Aderibigbe, and B.S. Afolabi, "Evaluating Web-Based Technologies: The Paradigm of User-Centricity," *J. Applied Computer Science and Mathematics*, vol. 10, no. 2, 2016, pp. 32–39.
7. S. Ray, "Essentials of Machine Learning Algorithms (with Python and R Codes)," Analytics Vidhya blog, 10 Aug. 2015; bit.ly/2xez1p3.
8. M. Iacob, *Good Practices in the Use of ICT in Providing Guidance and Counseling*, Inst. of Educational Sciences, 2012.
9. J. Bimrose and S.-A. Barnes, "Measuring the Effectiveness of Career Counseling," *Bildungsberatung im Dialog*, R. Arnold, W. Gieseke, and C. Zeuner, eds., Schneider Verlag Hohengehren, 2009, pp. 79–96.
10. J. Bimrose, J. Kettunen, and T. Goddard, "ICT—The New Frontier? Pushing the Boundaries of Careers Practice," *British J. Guidance and Counseling*, vol. 43, no. 1, 2015, pp. 8–23.
11. P.L. Benson, P.C. Scales, and M. Mannes, "Developmental Strengths and their Sources," *Handbook of Applied Developmental Science: Promoting Positive Child, Adolescent, and Family Development through Research, Policies, and Programs*, 2003, pp. 369–406.
12. J.P. Galassi and P. Akos, "Developmental Advocacy: Twenty-First Century School Counseling," *J. Counseling and Development*, vol. 82, no. 2, 2004, pp. 146–157.

Bernard Ijesunor Akhigbe *is a lecturer in the Department of Computer Science and Engineering at Obafemi Awolowo University, Nigeria. His research interests are in computing information systems generally, with a focus on information systems modeling and evaluation, end-user computing research, and human-computer interaction. Contact him at benplus1@gmail.com.*

Oyinkansola Onyinyechi Akhigbe *is a research student in the Department of Guidance and Counseling at the University of Ibadan, Nigeria. Her research interests include human development, guidance and counseling (G&C) using critical thinking models, and early childhood research. Contact her at akhigbeoyinkansola@gmail.com.*

Ishaya Peni Gambo *is a lecturer in the Department of Computer Science and Engineering at Obafemi Awolowo University, Nigeria. His research interests are in health informatics and software engineering with an emphasis on requirements engineering, architectural design, software quality improvement, and software testing. Contact him at ipgambo@gmail.com.*

Babajide Samuel Afolabi *is an associate professor in the Department of Computer Science and Engineering at Obafemi Awolowo University, Nigeria. His research interests are in computing information systems generally, with a focus in information systems modeling, information storage and retrieval, and user-oriented modeling. Contact him at afolabib@gmail.com.*

Adjusting to Autonomous Trucking

Shane Greenstein
Harvard Business School

News coverage of automation and machine learning tends to focus on extraordinary events, such as computers winning at Jeopardy and Go, and robotic arms flipping burgers in short-order restaurants. Additional headlines foster a sense of nightmares, conjuring pictures of autonomous cars killing pedestrians and newly automated establishments laying off their workforce. The combination of headlines has unleashed a near-hysteria, as if the near-future Terminator will either kill humans indiscriminately or rob the survivors of employment.

Let's come into contact with a grounded sense of the future. Some of this sour hype obscures the real technical improvement arising from advances in machine learning, neural networks, sensors, data access and retrieval, and systems engineering. Humans have invented tools for repetitive tasks, and some of those tools are becoming less expensive and more reliable.

This column stresses the basic economic lessons that inform a grounded view. Said simply, where the investments pay off quickly, firms will make use of tools. Experiments are looking for such payoffs right now. However, even with such payoffs, adjustments tend to take time and will slow adoption and change.

Autonomous trucking provides a good example. Considerable economic gains could be had from implementing some basic automation and machine learning in the trucking industry. Such automation will most likely be implemented at scale sometime in the next decade. What will happen to the 3.5 million truck drivers in the US who are focused on keeping their jobs? The basic point of this column is familiar to the experts in the area: Total employment is not at risk in the next decade, but it would not be surprising if a few job titles and assignments change. The hysteria simply gets it wrong.

MARGINS OF ADJUSTMENT

There is plenty of motivation for automating long-haul trucking. Of all trucking jobs, long haul is the most difficult to fill and staff. The work can be uncomfortable. It also creates enormous value in the US, annually carrying freight worth quadrillions—that is not a misprint (quadrillions comes after billions and trillions). Any productivity improvement yields enormous payoff.

Many of the increasingly useful and common forms of machine learning in trucking continue trends seen earlier. New algorithmic loops for error correction take advantage of improving processors. The training sets have become larger, and the maps include more detail, taking advantage of improving memory. Developers now see the potential for placing software on a server in one location, and it is possible to coordinate many users in other locations, taking advantage of improving network infrastructure.

 Published by the IEEE Computer Society

To be sure, trucking has already taken advantage of many advances in electronics. Most trucks contain on-board computers, GPS links, and numerous systems to monitor performance. In many trucks today, the software already moves with the machine. The sensors on mobile vehicles have also improved, and, again, the better software allows for a wider range of situations where the system can operate.

Not all is familiar, however. Many experts forecast that the next wave of sensors and software will address new applications in trucking, as part of a general wave of developments in autonomous vehicles. That forecast comes from borrowing progress in autonomous cars, and in spite of some basic technical issues that constrain progress in autonomous trucking. Trucks are much larger than most vehicles and carry a variety of payloads. Plus, they have their own patterns for accelerating and stopping. They also have rather different utilization records than a taxi or a commuter vehicle. The newest prototypes for autonomous trucking differ so much from those used in smaller cars that autonomous trucks cannot merely borrow the "training" from cars.

Trials in long-haul trucking involve training the vehicles for trips between depots adjacent to highways. At those depots, the trucks are handed off to drivers, who take them into cities for short-haul delivery. Judging from recent prototypes, humans are not disappearing anytime soon. Nobody is talking about installing robots in trucks to do the loading and unloading. The hard work today focuses on other high-value propositions, such as reducing safety issues from things like inattentive driving. A little automation can go a long way for that purpose—it can stop vehicles sooner, issue warnings to drivers, and relay information to dispatchers for use by others in a fleet. The prototypes also continue trends that began with the introduction of electronics into trucking long ago. Partial automation can enable longer continuous vehicle operation, better fuel consumption, and reduced maintenance expenses.

Judging from recent prototypes, humans are not disappearing anytime soon… Take the use of autopilot in commercial airlines today; software-enhanced navigation merely changes what the pilots do and when they apply their expertise.

So what limits progress? Like many applications in machine learning, there are too many "edge cases" that the software cannot yet satisfactorily handle—such as road construction, vehicles stopped at the side of the road, detours, pedestrians on the side of the highway, dead animal carcasses in the road, and so on. AI researchers know this problem well. Routine work is not as routine as it seems. Humans are pretty good at handling millions of variants of the little unexpected aspects of road work, police stops, bad weather, poor drivers, and breakdowns.

The statistics of edge cases are quite demanding. Software can be trained to handle much of this, perhaps 99 percent of the issues in a typical drive. But 99 percent is not anywhere near good enough. If, say, 1 percent is still left for humans, that translates into more than half a minute every hour in which a human needs to intervene. It is necessary to do much better than that to justify removing constant human awareness, and much better performance is required to get a sufficient return on the investment in the equipment to make it all work. In the lingo of the industry, partial or conditional automation is the most ambitious goal for the next several years. Full automation is a long way off.

SCALING

Any engineer in this area will say the same thing: There are many steps between prototypes and large-scale implementation of a fleet of partially autonomous vehicles. Now I say this: Scaling is exactly the topic in which some economic reasoning provides better guidance than hysterical forecasts.

This article originally appeared in IEEE Micro, *vol. 38, no. 3, 2018.*

Scaling requires predictable business processes that can be measured and monitored. The drivers might actively drive less, but they still might help with aspects that affect fueling, safety, liability, and loading and unloading. Take the use of autopilot in commercial airlines today; software-enhanced navigation merely changes what the pilots do and when they apply their expertise.

We should expect that business processes will adjust and adopt new routines. The timing for fueling, maintenance, docking, and inspection will change. New procedures for monitoring daily, weekly, and monthly targets will be put in place. Will that eliminate work? No, but it will shift who does the work, what they do, and how they are trained to do it.

The new timing will require new principles for organizing teams. The new teams will require new principles for responsibility. For example, who pays for the costs of error—the driver or the programmer? Beyond that, economics provides reasons to be cautious. Learning will lead to new services—say, driving in the middle of the night when there's less traffic. Those new services might also come with new potential logistical limitations, such as inability to drive in certain weather or road conditions.

Will that change work at the organizational level? Yes, sure, and even if we cannot precisely forecast how, it is clear where the type of work will change. The planning department will tackle new issues. So will the legal team, the logistics department, the business partnership liaisons, the sales department, and the billing group. That will feed back into the R&D department, which will be asked to change what patents it applies for. Trucking companies will adjust on many margins to accommodate autonomous trucking at scale.

And here is the kicker. The reduction in cost might generate *more* demand for services, which might lead to *more* employment of truckers. It is hard to forecast the totality of all this change.

LESSONS

There was a point to this thought exercise. Let's review the broad lessons.

There will be adjustments. Tasks will change and so will the regular processes that accompany the execution of these tasks. Team assignments and composition of teams will change, and so too will organizations that manage partnerships with these teams. These changes will be just as difficult to make as the technical inventions that precipitated them.

It is easy to see how the software will improve and bring about costs savings. Most of the consequences are, however, rather unpredictable. When will the biggest technical gains emerge? Will the cost savings be substantial enough to alter productivity and pricing?

And what about employment? Will total employment in trucking go up or down in the next decade as a result of the increasing use of autonomous vehicles? While tasks might change in ways to diminish the need for some work, the costs might lead to an increase in volumes and increase the demand for work. On net, there is no way to predict with certainty. That said, massive layoffs or other employment nightmares are highly unlikely.

CONCLUSION

Gains in neural networks have gone beyond what had been widely appreciated in public conversation. There is much productivity improvement on the near-term horizon. More to the point, there is not much merit to most of the hysteria behind machine learning. There is a lot we have to learn. In the meantime, we all need to keep on trucking, because society is all in for the long haul.

ABOUT THE AUTHOR

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

IEEE COMPUTER SOCIETY

2018 CS Election

Volunteer Leadership is Vital

Voting for IEEE Computer Society candidates for president-elect, first and second vice presidents, and Board of Governors members is open through 24 September, 2018.

Visit **www.computer.org/web/election** for information and to cast your ballot!

CONNECT ON INTERFACE

Explore **INTERFACE**, a communication resource to help members engage, collaborate and stay current on Computer Society activities. Use **INTERFACE** to learn about member accomplishments and find out how your peers are changing the world with technology.

We spotlight our professional sections and student branch chapters, sharing their recent activities and giving leaders a window into how chapters around the globe grow, thrive and meet member expectations. Plus, **INTERFACE** will keep you informed on Computer Society–related activities so you never miss a meeting, career development opportunity or important industry update.

Connect today at interface.computer.org

IEEE COMPUTER SOCIETY
INTERFACE

CALL FOR PAPERS

Authors are invited to submit manuscripts that present original unpublished research in all areas of parallel and distributed processing, including the development of experimental or commercial systems. Work focusing on emerging technologies and interdisciplinary work covering multiple IPDPS areas are especially welcome. Topic areas include:

- **Parallel and distributed computing theory and algorithms (Algorithms):** Design and analysis of novel numerical and combinatorial parallel algorithms; protocols for resource management; communication and synchronization on parallel and distributed systems; parallel algorithms handling power, mobility, and resilience.
- **Experiments and practice in parallel and distributed computing (Experiments):** Design and experimental evaluation of applications of parallel and distributed computing in simulation and analysis; experiments on the use of novel commercial or research architectures, accelerators, neuromorphic architectures, and other non-traditional systems; algorithms for cloud computing; domain-specific parallel and distributed algorithms; performance modeling and analysis of parallel and distributed algorithms.
- **Programming models, compilers and runtimes for parallel applications and systems (Programming Models):** Parallel programming paradigms, models and languages; compilers, runtime systems, programming environments and tools for the support of parallel programming; parallel software development and productivity.
- **System software and middleware for parallel and distributed systems (System Software):** System software support for scientific workflows; storage and I/O systems; system software for resource management, job scheduling, and energy-efficiency; frameworks targeting cloud and distributed systems; system software support for accelerators and heterogeneous HPC computing systems; interactions between the OS, runtime, compiler, middleware, and tools; system software support for fault tolerance and resilience; containers and virtual machines; system software supporting data management, scalable data analytics, machine learning, and deep learning; specialized operating systems and runtime systems for high performance computing and exascale systems; system software for future novel computing platforms including quantum, neuromorphic, and bio-inspired computing.
- **Architecture:** Architectures for instruction-level and thread-level parallelism; memory technologies and hierarchies; exascale system designs; data center architectures; novel big data architectures; special-purpose architectures and accelerators; network and interconnect architectures; parallel I/O and storage systems; power-efficient and green computing systems; resilience and dependable architectures; performance modeling and evaluation.
- **Multidisciplinary:** Papers that cross the boundaries of the previous tracks are encouraged and can be submitted to the multidisciplinary track. During submission of multidisciplinary papers, authors should indicate their subject areas that can come from any area. Contributions should either target two or more core areas of parallel and distributed computing where the whole is larger than sum of its components, or advance the use of parallel and distributed computing in other areas of science and engineering.

The five-day IPDPS program includes three days of contributed papers, invited speakers, industry participation, and student programs, framed by two days of workshops that complement and broaden the main program.

GENERAL CHAIR

Vinod Rebello *(Fluminense Federal University, Brazil)*

PROGRAM CHAIR AND VICE-CHAIR

José Moreira *(IBM Research, USA)* and
Alba Cristina Melo *(University of Brasilia, Brazil)*

PROGRAM AREA CHAIRS AND VICE-CHAIRS

• ALGORITHMS:
Gianfranco Bilardi *(University of Padova, Italy)* and
Denis Trystram *(Grenoble Institute of Technology, France)*

• EXPERIMENTS:
María Jesús Garzarán *(Intel Corporation, USA)* and
Saeed Maleki *(Microsoft Corporation, USA)*

• PROGRAMMING MODELS:
Xavier Martorell *(Technical University of Catalunya, Spain)* and
Christian Terboven *(RWTH Aachen University, Germany)*

• SYSTEM SOFTWARE:
Dilma Da Silva *(Texas A&M University, USA)* and
P. (Saday) Sadayappan *(Ohio State University, USA)*

• ARCHITECTURE:
Per Stenström *(Chalmers University of Technology, Sweden)* and
Rodolfo Azevedo *(University of Campinas, Brazil)*

• MULTIDISCIPLINARY:
Nancy Amato *(Texas A&M University, USA)* and
Andrea Pietracaprina *(University of Padova, Italy)*

Abstracts due	**October 8, 2018**
Full paper due	**October 15, 2018**
Preliminary decisions	**December 10, 2018**
Final submissions due	**January 7, 2019**
Final notification	**January 21, 2019**

IPDPS 2019 VENUE

Rio de Janeiro is known for its natural settings, plentiful beaches, dramatic mountains, and a vibrant city center that resonates with culture and a deep sense of history and heritage. Its wonderful weather, unique gastronomy, and famous landmark attractions draw visitors from around the world. Join IPDPS 2019 in Rio at the Hilton Rio de Janeiro Copacabana to experience first-hand why this is such a special and memorable destination!

IEEE COMPUTER SOCIETY
Sponsored by IEEE Computer Society Technical Committee on Parallel Processing

In cooperation with ACM SIGARCH & SIGHPC and IEEE TCCA & TCDP