Carnegie Mellon University
Software Engineering Institute

EMERGING TECHNOLOGIES 2020: SIX AREAS OF OPPORTUNITY

Charles Holland
Jacob Tanenbaum
December, 2020

Table of Contents

Introduction

The purpose of this study is to understand what the software engineering community perceives to be key emerging technologies. In this context, “emerging” describes a technology that has yet to mature but is undergoing a period of rapid innovation and change. Our perspective focused on technologies themselves (innovations in hardware and software) rather than related factors (such as processes and policies).

The six emerging technologies described here hold great promise and, in some cases, have already attracted the interest of the Department of Defense (DoD). By understanding these technologies and their intersection with DoD needs, we can create a research agenda that serves our sponsor’s mission.

Constraints, Limitations, and Assumptions

We conducted this study using unclassified and publicly available sources. The references list presented here is by no means exhaustive, but rather highlights six key technological areas. Our discussion of these areas focuses on novel topics, rather than familiar and well-documented topics. While we seek to provide an introduction to these technologies, the resources cited in the reference list can provide much deeper insights into the potential and capabilities of these technologies. Because these technologies are constantly changing, this document is a snapshot of what is known of the current landscape. It is the product of work done over the course of only a few months.

Acknowledgements

The authors benefited from substantive interactions with colleagues across the SEI, Carnegie Mellon University (CMU) communities, and others in the software engineering community during the preparation of this report. The authors would like to specifically thank Matt Gaston, John Wohlbier, Greg Such, and Ben Cohen.

About the Sources

Our goal was to understand the perspectives on emerging technology from all stakeholders in contemporary software engineering. We consulted publications from across numerous domains, including defense, information technology, consulting, and academia. These reports often included explicit lists of key technological themes of significance to the organization. Many were intended for larger audiences than the software engineering community, and included themes that are out of scope or tangential to software engineering.

As a federally funded research and development center (FFRDC), the opinions of the defense sector are essential to SEI. Consequently, we found the following NATO and DoD publications particularly significant. The North Atlantic Treaty Organization (NATO) produced a 140-page report entitled *Science & Technology Trends 2020-2040: Exploring the S&T Edge*. This lengthy document classified technologies into four themes: intelligent (having to do with autonomy and decisions), interconnected (related to networks), distributed (deployed in many distant places), and digital (fusing physical and information domains). It featured lengthy sections on several key technologies, including artificial intelligence (AI) and quantum computing.

The U.S. Department of Defense's *Digital Modernization Strategy* listed technologies that were key to four organizing goals and their four CIO priorities: cybersecurity, AI, cloud, and command, control and communications (C3). The document outlined thirteen objectives, each composed of its own elements, including the establishment of the Joint Artificial Intelligence Center (JAIC), treating data as a strategic asset, and migrating systems to the cloud. In the section "Technologies Offering Promise to DoD," it listed key technologies, some of which were emerging and others that were more mature.

The consulting industry and related sources also provided several useful lists of resources. Gartner, Inc., a prominent player in market forecasting of technology, is notable for its "hype cycle" model, which postulates five phases of evolution for a technology. The cycle starts when a technological breakthrough generates sudden enthusiasm. The technology later fails to meet expectations and eventually becomes fairly understood and useful. Gartner provides deep market research on a variety of specific technical areas, which means it is familiar with many details in a breadth of technical sectors. Gartner's emerging technologies list also described human factors relating to the enumerated themes.

In its report, "IDC FutureScape: Worldwide IT Industry 2020 Predictions," the International Data Corporation presents a series of predictions of technological adoption rates, each using a specific year before 2026 as a milestone, and describing various technological landscapes. Forbes' "The 7 Biggest Technology Trends In 2020 Everyone Must Get Ready For Now" lists seven technologies that will become relevant from a business perspective. Accenture's Tech "Vision 2020: Technology Trends to Watch" examines just five themes more deeply, looking at their implications in various domains, elaborating on aspects of each theme, and detailing why each one is relevant to business. They describe the concept of the "beta burden," which describes the perpetual mutation of modern technology—to even include more mature technologies—putting it in a "beta" state of instability. Similarly, Deloitte's substantial "Tech Trends 2020" report highlights the theme of "Architecture Awakens," or the importance of integrating large-scale system design to the organizational enterprise. The report featured sections describing six themes, including case studies and executive perspectives.

The research community, as the birthplace of so many new technologies, is an essential perspective for assessing emerging technologies. The Computing Community Consortium (CCC), a group supported by the National Science Foundation and focused on empowering innovative, high-impact research, lists eight key categories pertaining to emerging technologies in computer science. Notably, it includes an explicit category for innovation in theory and algorithms. Each of the CCC computer science areas has its own set of CCC publications, including white papers, workshop reports, etc. Massachusetts Institute of Technology's article "10 Breakthrough Technologies 2020" discusses ten novel topics across numerous fields unmentioned in other articles. The technologies listed were more specific, such as satellite mega-constellations and AI-discovered molecules. Each topic includes an explanation of its significance, availability estimates, and key players.

The Networking and Information Technology Research and Development (NITRD) Program constitutes expert government perspectives in advanced computing and IT. Its "Program Component Areas" list of key technologies informs a supplement to the presidential budget. The list includes research and development of emerging technologies, such as computing-enabled networked physical systems and high-capability computing systems, as well as present-day concerns, such as education and workforce. Of the Institute of Electrical and Electronics Engineers (IEEE) Computer Society's list of top twelve technology trends, half were related to machine intelligence: "AI at the edge," "AI and critical systems," "cognitive skills for robots," "AI/ML applied to cybersecurity," adversarial machine learning, and "reliability and safety challenges for intelligent systems." IEEE Computer dedicated its December, 2019 issue to discussing many of these themes.

Other sources we consulted did not fall into the previous categories. Industry Week's "Top 10 Technologies to Watch in 2020" briefly lists emerging technologies from a manufacturing perspective, specifically highlighting the industrial Internet of Things (IoT) as well as 3D printing and additive manufacturing. The World Economic Forum's "Top 10 Emerging Technologies 2019" report examined emerging technology from a broad, holistic perspective, encompassing the fields of medicine, agriculture and energy. Y Combinator, a Silicon Valley organization that provides seed funding for startups, produces a *Requests for Startups* guide that lists key topics for innovation in the near future. Its 21-item list also addresses many disciplines in addition to software, with many topics more germane to policy than technology, but provides an otherwise-lacking startup financing perspective. Finally, we also studied the opinions of SEI business developers, as transcribed from SEI's *National Agenda for Software Workshop 2020* event. Our colleagues in this role are in constant communication with DoD and federal government personnel, and have an essential perspective on what priorities are on the minds of these key stakeholders. The business developers' perspectives were often more focused on the near term, and they spoke precisely to customer's requirements.

Other, nonconventional sources we did not examine might hold potential. For instance, we did not examine science fiction. While less concrete even than research, science fiction offers novel visions about the future and might be relevant to what actually emerges in the future. The 1983 film *War Games* awakened President Reagan to the significance of the emerging cybersecurity field. The aforementioned NATO report referenced a "Strategy of Technology" written by science fiction author Jerry Pournelle. DARPA is often interested in science fiction works in general.

Approach

The SEI team developed this list of six technological themes based on a survey of the available literature noted in the previous section. To make its selections, the team analyzed the list through the lenses of 1) level of technical interest, and 2) opportunity for DoD. In addition, significant opportunities for combining multiple technologies exist to multiply capability. These opportunities present substantial challenges for software engineering.

Technology Theme	MIT Tech. Rev.	NATO	Deloitte	Forbes	Gartner	Industry Week	IDC	Accenture	World Economic Forum	IEEE	Y Combinator	SEI BD	DoD Digital Modern Strat.
Advanced Computing	Yes					Yes				Yes		Yes	Yes
The Smarter Edge	Yes	Yes			Yes	Yes	Yes			Yes			Yes
Digital Twins		Yes	Yes		Yes					Yes			
AI		Yes		Yes	Yes	Yes	Yes			Yes	Yes	Yes	Yes
Extended Reality		Yes	Yes	Yes	Yes	Yes		Yes	Yes		Yes		
Data Privacy, Trust, and Ethics	Yes	Yes	Yes	Yes	Yes	Yes	Yes		Yes	Yes			Yes

Table 1: Matrix of Prevalence of the Six Technological Themes in Publications.

Advanced Computing

We use the term *advanced computing* to refer to advances in computing power that are enabled by advances in computer hardware, from the chip level to the system level. Advanced computing is the driver for new capabilities that are enabled through software. In addition to traditional semiconductor-enabled computing, quantum computing is an important emerging technology that will revolutionize computing.

Semiconductor advances have enabled Moore's Law—the doubling of transistors' power with every generation. The decreasing size of transistors is accompanied by Dennard's voltage scaling, which described more than a doubling in chip computing power (performance) for the same amount of electrical power. Dennard's scaling ended roughly in 2005, and the chip industry has shifted over the last decade to multicore, graphics processing units (GPUs), field programmable gate arrays (FPGAs), and application-specific integrated circuits (ASICs) at the chip level for continuing the trajectory of improved performance. All of these shifts require advances in software engineering.

What is especially new is the attention to developing a hardware strategy to satisfy requirements for machine learning, both for large problems that require supercomputing and much smaller chips at the edge for inference. This area presents a very diverse problem set. Historically, maintenance of the stockpile, cryptography, and challenging scientific problems (such as weather prediction and climate change) drove federal investment in DoD applications. Today, we are seeing a drive for applications to support COVID-19 vaccine design.

Quantum computing has the potential to revolutionize approaches to these and other problems. However, we are in the early stages of that journey. The recent push for exascale computing is being led by the United States Department of Energy (DoE) for science and Microsoft and others for AI. In 2015, the National Strategic Computing Initiative (NSCI) started exascale computing focused on traditional supercomputing plus big data challenges. Exascale machines are arriving. In August 2019, the NSCI Fast Track Action Committee provided an update with a broader vision: "Pioneering the Future of Computing."

Subfields of Interest

AI is driving supercomputing and vice versa. NVIDIA CEO Jensen Huang uses the phrase "Cambrian Explosion" to describe innovation in neural network algorithms and the specialized hardware for implementing them [Freund 2019]. Cerebras has developed the Wafer Scale Engine (WSE) that boasts "1.2 trillion transistors, 400,000 processor cores, 18 gigabytes of SRAM, and interconnects capable of moving 100 million billion bits per second" [Moore 2020]. The WSE is designed to enable rapid training of large neural networks.

Microsoft has invested in supercomputing for AI with its Massive AI Supercomputer on Azure. The system features 285,000 CPU Cores and 10,000 GPUs. Microsoft created it "for training larger AI models targeting highly complex problems" [Black 2020].

Advances in AI will require new software to run on these systems, opening up new opportunities for software engineering.

Today, information can be encrypted only for transmission and storage. Fully homomorphic encryption (FHE) makes it possible to analyze or manipulate encrypted data without revealing the data to anyone, a major advance. FHE builds upon Craig Gentry's seminal 2009 work and other work to date, initially a million times too slow to be practical [Gentry 2009]. A new DARPA MTO program, Data Protection in Virtualized Environments (DPRIVE) for FHE, seeks to design and implement a "word size" (i.e., the core unit of data for the CPU) of 1000-bits for its hardware accelerator. This would reduce the computational runtime of FHE algorithms to be only 10 times slower than the unencrypted default. IBM has already released a fully homomorphic encryption toolkit for Mac OS and IOS, and its Linux and Android toolkits are on the way.

Small-scale quantum computers that use various technologies for qubits are emerging. These development efforts are led by major companies (IBM, Honeywell, Google, Microsoft, etc.) and venture capital-funded activities. Quantum computing resources are now available on the cloud. While available quantum resources range to up to a few hundred qubits, they are insufficient for error correction. This technology level is termed Noisy Intermediate Scale Quantum (NISQ) due to its characteristic unreliability. The challenge is to demonstrate commercial and/or economic benefit with NISQ machines. This will enable a virtuous cycle, similar to semiconductor technologies over the past 40 years. It will also demonstrate the advantage of applying quantum computing to valuable problems. A longer-term technological opportunity exists to develop a software ecosystem that enables scalable quantum computing.

Opportunities

Quantum computers, upon achieving the necessary scale, could be used to break contemporary public key cryptography (PKI). This type of encryption is a bedrock of information security. Breaking it would cause a major disruption of secure enterprises everywhere. Today's best estimate on algorithm requirements can be found in Gidney and Ekera's 2019 paper "How to factor 2048-bit RSA integers in 8 hours using 20 million noisy qubits" [Gidney 2019]. This new capability to threaten RSA places the entire PKI system at future risk. The Federal government has tasked the National Institute of Standards and Technology (NIST) to lead a major effort to decide on the quantum resistant algorithms that will replace PKI as the new standard for encryption.

A major overhaul of encryption will be needed after these new cryptographic algorithms are developed. The process of crypto-modernization will be a substantial, decade-long event; its implementation will be a software engineering opportunity. Information systems in all domains will need to be re-engineered to support the new standard. Thorough verification and validation will be mandatory to assure system security.

The Smarter Edge

The Smarter Edge is a catch-all term for new developments in a system of heterogeneous computing devices. The system goes beyond a conventional computer network and incorporates devices at the edge of the network, including sensors, Internet of Things (IoT) devices, and mobile phones. The Smarter Edge serves as an alternative architecture to cloud computing because it shifts much of the computing burden to local edge devices instead of centralized data centers.

While the concept of ubiquitous computing has existed for decades, there have been recent advancements to accelerate the Smarter Edge. Improvements in computing hardware enable more powerful programs to be executed at the edge. The arrival of 5G networks raises the bandwidth ceiling. Finally, data scientists are transforming and adapting machine learning (ML) algorithms to be better suited for smaller devices, developing a "Tiny AI." Mark Weiser (Xerox PARC) famously detailed this concept in his 1991 article, "The Computer for the 21st Century" [Weiser 1991].

Edge data is growing, thanks to new sources such as ubiquitous sensing and the IoT. Improvements in computer hardware are enabling the development of more complex, advanced software. Computing resources may be organized in new structures, such as fog computing or cloudlets. The field of analytics is creating innovative new ways to examine data. In AI, algorithmic improvements allow a smaller resource footprint.

Subfields of Interest

The new practice of Tiny AI explores the miniaturization of AI and ML. Edge devices are typically much lighter weight than even a single cloud server. Moving AI to these devices imposes many constraints, including ultra-low power, a small resource footprint, and minimal library and/or binary dependencies. The inaugural 2019 TinyML Summit attracted more than 90 companies and included Karl Pfister, the originator of Smart Dust in the 1990s, as a speaker. Qualcomm is conducting significant research in the field. The emerging interdisciplinary field of neuromorphic computing seeks to deeply model brain-like activity (far beyond traditional neural networks) and offers the potential to deliver AI functionality for much lower energy costs.

The advent of 5G networks will ramp up the available bandwidth for IoT devices and the Smarter Edge in general. The DoD 2019 Modernization Strategy terms 5G a "technology offering promise," citing improvements in latency, cell edge performance, and efficient use of the spectrum. However, there are some risks associated with 5G. Firms closely tied with China and other countries have dominant presences in component manufacturing. This risks introducing untrusted hardware and/or software to the enterprise. In addition, to support backward compatibility, 5G may inherit vulnerabilities from legacy networks.

Beyond stationary IoT devices, drones could be integrated into the Smarter Edge as drone swarms. Drones would be a more complex component of the smarter edge, facing extra connectivity challenges due to their mobility. The number of drones (perhaps ranging from tens to hundreds) in the swarm might also introduce scaling challenges.

The future Smarter Edge may include more nontraditional devices. Space-based systems such as satellite mega-constellations could enable the next wave of connectivity. This scenario is enabled by the emerging potential for lower-cost constellations of small satellites. SpaceX recently announced a launch of 32,000 Linux computers to enable Starlink Internet [Tung 2020]. DARPA's Blackjack concept builds upon advances in private sector, low Earth orbit (LEO) satellites to support military payloads and spacecraft [Forbes 2020]. A system of several hundred LEO-based satellites, while challenging to implement, could empower hypersonic cruise missile defense. The DoD and the commercial sector have many other similar endeavors.

Opportunities

The Smarter Edge offers opportunities to improve many system qualities, including

- bandwidth
- latency
- outages
- security
- privacy
- power awareness

Major application drivers today include the health, manufacturing, predictive maintenance, and autonomous domains.

Digital Twins

A *digital twin* is a high-fidelity digital or computer representation of a physical object with some ability to reason about its properties. Originally, those properties were related to data with respect to the structure of the object (e.g., CAD drawing or assembly properties). They did not represent the behavior of the object. However, digital twins have expanded to include properties that represent the function or performance of the object in the real world. This model empowers experimentation and prediction. Digital twins are an evolving concept that has gained much traction in the community.

The most recent trend for digital twins is the transmission of real-time data sensed by the actual object back to its digital twin. This new, higher-resolution sensor data allows the digital twin to reason about the object's future behavior, then transmit feedback to the physical object. In turn, the physical object could use this feedback from its digital twin to inform tasks such as prediction and control messaging.

One framework posits four stages of the evolution of digital twins. The most advanced digital twins today are in stage three, where the digital and physical twin are able to synchronize data and status. The final stage is conceived as a convergence of the digital and physical twin, where the digital twin is able to duplicate itself and interconnect with other digital twins while still maintaining its identity and relationship with the physical twin.

The expansion of the digital twin concept is also driven by a more complete technological ecosystem, taking advantage of advanced computing, visualization capabilities, real-time sensor data, etc. Electronic prototyping is another component of this concept, enabling reasoning about the digital object before the actual physical object is built. Current virtual prototyping efforts are being driven by advances in high-performance computing (HPC) capability and advances in scientific computing algorithms representing complex physics. One example of this work is the DoD's High-Performance Computing Modernization Program (HPCMP), including the Computational Research and Engineering Acquisition Tools and Environments (CREATE) program recently led by the SEI's Doug Post.

Digital twins have broad applications because the actual object under consideration could vary greatly. Circa 2018, the Singapore National Research Foundation produced Virtual Singapore, a digital model of an entire city that was used for planning but not for real-time feedback [Singapore 2020]. Another growing area is the use of digital twins in enterprise-wide business operations and manufacturing. IBM predicts that digital twins in farming will produce efficiency windfalls, greatly increasing agricultural yields [IBM 2019].

Opportunities

The Undersecretary of Defense for Research and Engineering (USD(R&E)) recognized the importance of digital twins in the June, 2018 DoD Digital Engineering Strategy [DoD 2018]. Also, NATO believes that digital twins will improve logistical efficiency and reduce life-cycle costs [Reding 2020].

The interaction and cooperation of multiple digital twins has not yet been deeply explored.

Another important question for research is how to quantify uncertainty. A digital twin is never a precise replica of reality. How can we quantify the uncertainty in this model? One approach is to integrate AI with digital twins. Digital twins can provide a deterministic, systemic view of an entity, while AI, given sufficient training data, can provide predictions for cases where the entity is not completely understood. AI may also be able to effectively employ newly-available real time feedback data.

While digital twin modeling has focused on mechanical processes thus far, it has great potential in modeling living entities. Projecting a human into a digital twin could produce richer models. Modeling physiology and biological processes could revolutionize medicine. For example, digital twin modeling could provide doctors with interactive environments to prepare for surgery or prescription.

Any digital twin system will need to consider data security. Protecting the data of individuals and enterprises is increasingly important for the successful IT ecosystem. Given the data-centric nature of digital twins, this concern is only amplified.

Artificial Intelligence

Artificial intelligence (AI) is now common across industries and in the public consciousness. Its ubiquity underwrites its importance. AI is a field unto itself that encompasses numerous subfields, each of which is seeing constant growth and research.

AI is the phenomenon of machines making complex decisions. Within AI, the subfield of machine learning (ML) is currently dominant. ML programs train from experience, or additional data, heavily relying on statistics to make predictions. One class of algorithm, deep learning, models reality using a neural network—a series of interconnected neuron nodes, resulting in a complex system of equations. Given large data sets and powerful computers, deep learning often produces the most accurate predictions available today.

Subfields of Interest

Adversarial machine learning is the harmful manipulation of an AI system by an attacker. In interviews with 28 organizations, Microsoft found widespread deficiencies of protection from attacks on organizations' machine learning systems [Kumar 2020]. Common attacks include exploratory attacks (to understand details of data and system parameters), evasion attacks (to force the system to produce an undesired prediction), and poisoning attacks (giving the system bad data to corrupt it in the future). Researchers are continuously developing new exploitations and many papers use image processing as examples. The National Institute of Standards and Technology (NIST) developed a taxonomy to classify types of adversarial attacks. Advances in quantum computing may be accompanied by new challenges from quantum adversarial machine learning [Lu 2020].

While deep learning provides valuable predictions to decision-makers, these predictions are typically lacking in explanation. Even academic specialists struggle to understand why a deep neural network makes specific predictions. The emerging field of causal learning seeks to transform this issue. Causal learning analyzes a problem, searching for notably strong relationships between variables that imply a cause-and-effect relationship between them. Unless a complete, human-level intelligence is developed, humans will need to be involved in making decisions. Human-AI Interaction analyzes the relationship between human users and AI-enabled systems. At Microsoft, Amerishi et al. compiled "Guidelines for Human-AI Interaction," which included recommendations like supporting efficient dismissal and correction and to "Make clear how well the system can do what it can do" [Amershi et al. 2019].

Interests

All kinds of organizations around the world are expressing interest in AI. The European Union (EU) is investing in AI through Horizon 2020 and its predecessor programs (e.g., RAISE, the First International Workshop on Realizing Artificial Intelligence Synergies in Software Engineering). China's "New Generation Artificial Intelligence Development Plan" diverted billions of dollars into AI research and development [China 2017]. Analysis by the Allen Institute for Artificial Intelligence in Seattle showed that the proportion of Chinese authorship in the top 10% most-cited papers is rising [O'Meara 2019]. There is also Russian interest at top government leadership levels.

In the academic community, the Computing Community Consortium (CCC) produced *A 20 Year Community Roadmap for Artificial Intelligence Research in the US* led by Yolanda Gil (President of AAAI) and Bart Selman (President Elect of AAAI) [Gil 2019]. They are building a consensus around research visions and creating funding opportunities to enable them.

Major technology industry leaders are committed to AI, include Microsoft, Google, Facebook Amazon, Qualcomm, and Intel. From Google AI's Jeff Dean's perspective, "We want to use AI to augment the abilities of people, to enable us to accomplish more and to allow us to spend more time on our creative endeavors" [Dean 2020]. Peter Lee is leading Microsoft AI for Health, a major new initiative to apply AI to healthcare.

The DoD has great interest in AI technology. The Joint Artificial Intelligence Center (JAIC) is the focal point for the DoD's utilization of AI. In September, 2020, JAIC hosted the Department of Defense Artificial Intelligence Symposium and Exposition, a two-day workshop. JAIC recently produced "Understanding AI Technology," an overview of AI and ML technology designed for non-technical managers, officers, and executives.

Autonomous vehicles are another major source of investment by businesses that range from the automotive industry to the IT industry. Cities across the country are home to prototype vehicles driving on local road networks, learning through real-world experience to become more autonomous. The DoD is interested in autonomous vehicles of its own, but faces greater challenges. The Army cited issues that included the following [Army 2020]:

- much lower volumes of training data available
- less time available to train autonomous vehicles
- driving environments that are off-road and rugged
- generally adversarial and hostile conditions

Opportunities

AI has the potential to enable better software engineering. Example applications of AI technology include better searching and completion of code, auto-repairing of programs, and bug detection. Many organizations are pursuing this topic. OpenAI is a renowned research lab dedicated to AI studies. It recently began offering its powerful text generator as a service. The newest iteration, GPT-3, is even more capable than ever. It is able to generate simple code (e.g., UI) based on human language descriptions [Vincent 2020]. Microsoft is offering its Deep Coder and Code Defect tools. Amazon provides the Code Guru suite. This application includes a profiler tool to help developers find an application's most expensive lines of code and features specific visualizations and recommendations on how to improve code to save money. It also includes a reviewer tool to improve code quality, which uses machine learning to identify critical issues and hard-to-find bugs during application development.

Just as software engineering can benefit from AI, AI can benefit from software engineering. Researchers are analyzing how to improve AI by viewing it as a software system [Sculley et al. 2015]. A structural challenge to AI engineering is the non-determinism of AI. Unlike a traditional application, stakeholders may not be able to definitively state what an AI system "should" say given a set of user input, making testing much more challenging. Google developed an ML Test Score that decomposes an AI system and assesses it at each element of its pipeline (e.g., features and/or data, model development, and infrastructure).

AI will continue to be used as a service in which customers grant the AI system access to their data and receive a prediction or recommendation in return. This will enable new business models. For example, an AI system could recommend predictive maintenance for jet engines, enabling the sale of jet engines with costs by the hour of use. AI also has potential as an automated assistant for humans. Decision-making is especially important in the intelligence community for rapidly developing and providing documentation, especially in crisis situations.

From CMU Professor Martial Hebert's perspective, we need to "look at AI very broadly, from the physical layers (sensors), to software, to the ML algorithms, to human interactions and the social sciences." [Carleton et al. 2020]. Today, many AI practitioners approach prediction problems from a generalist perspective, lacking context about the problem domain and its data. Empowering data scientists with knowledge in both AI and domain data offers the potential to produce AI systems of higher quality and relevance. In general, AI has the potential to democratize many domains by lowering the skill (and money) required to execute important tasks.

Extended Reality

Extended reality refers both to augmented reality (AR), virtual reality (VR), and combinations of the two. AR refers to the use of devices, such as specialized glasses, that display additional information about the scene of view. Here the individual sees the real scene but with augmented information. VR, in contrast, refers to wearing specialized devices where the human only sees the virtual world. An essential quality of extended reality is its power to radically reshape humanity's reasoning about information. Overall, extended reality will transform the way we work, build, create, and collaborate.

Extended reality has historical roots in research and academia. The Electronic Visualization Laboratory (EVL) at the University of Illinois at Chicago developed CAVE Automatic Virtual Environment (CAVE–a recursive acronym). CAVE is a science-based facility for visualizing supercomputing data. Circa 2007, DARPA developed Deep Green, a real-time computation of course of action, which was projected onto wearable glasses. In the past few years, advances in graphics processing units (GPUs) have commoditized this field and enabled a potential revolution. Although entertainment is the recent driver for this technology, many expect to see enterprise use increasing.

Several emerging technologies enable and empower extended reality. 5G networks will create the spaces in which extended reality can function, such as conducting a VR–AR meeting from a taxi. Cloud computing is another natural complement. It provides powerful processing of data offloaded by peripheral devices, possibly routed through a mobile phone as a middleman. When latency is a concern, edge computing and Tiny AI can be substituted for the cloud.

AI and natural language processing technologies have been used in recent years to create basic machine translators of human language. Companies are building on these technologies to develop real-time translation that is executed in the cloud, effectively embedding this capability into headphones. This makes possible smooth conversations between speakers without a common language. Wavery Labs, iFlytex, and Pocketalk are examples of vendors developing real-time translation headsets.

The private sector sees a lot of promise in extended reality. Based on its analysis, Deloitte observed "companies focusing on the human experience have been twice as likely to outperform their peers in revenue growth over a three-year period, with 17 times faster revenue growth than those who do not." [Deloitte 2020]. Similarly, Accenture found that "worldwide spending on AR and VR is expected to pass $18 billion in 2020—a 78.5% increase over 2019—and will reach $160 billion by 2023" [Daugherty et al. 2020].

Major technology players are investing in their own extended reality offerings. Amazon is offering its Sumerian tool to enable AR/VR deployments. Sumerian interfaces with Amazon Web Services (AWS) and produces displays that are portable to common AR/VR platforms. Similarly, Apple is offering its Swift Playgrounds tool and demonstrating its capabilities through its [AR]T project. This project features AR displays in major cities of interactive art created for the platform by prominent contemporary artists.

In May 2020, Spatial announced that its eponymous virtual collaboration application would become free. The Spatial tool renders an environment (e.g., a conference room) in the cloud. Users can sign in and jointly work and communicate in a virtual, three-dimensional space. Given the trend of working from home triggered by COVID-19, these virtual collaboration tools may become popular productivity enablers.

Opportunities

Overall, extended reality has drivers in almost every domain. Video games and entertainment are intuitive applications. There is also potential in sectors such as health care, real estate, military, science, and education. VR provides realistic training that simulates scenarios and environments that would be unthinkably expensive or risky to reproduce otherwise.

Extended reality can also incorporate diverse human behaviors and experiences into simulations at reduced cost and complexity. This is especially important for DoD and government applications. Specifically, there are large DoD opportunities in training and simulation, diagnostic repair, and operations. The Army is experimenting with this technology in operations in the field. The DoD focus on modeling and simulation for realistic and efficient operational training is known as Team Orlando. It encompasses all DoD service branches plus the United States Department of Homeland Security (DHS), with members co-located in Orlando, FL. The organization boasts that "members of Team Orlando touch simulation and training systems that are used by virtually every soldier, sailor, marine and airman" [Team Orlando 2020]. Similarly, NASA is also using extended reality to plan future Mars projects.

Extended reality offers many areas for innovation. The Carnegie Mellon University Future Interfaces Group, led by Professor Chris Harrison, works beyond traditional VR and AR human-computer interfaces. This group is developing novel sensing and interactive technologies that couple with ML. Researchers are also studying devices and methods to unobtrusively, neuroscientifically measure thought. Examples include electroencephalography (EEG), eye tracking, facial coding, galvanic skin response, and implicit association testing. Open questions remain regarding how to improve interoperability among different AR/VR devices and the scalability of extended reality platforms.

Data Privacy, Trust, and Ethics

Data is now a strategic asset. It is bundled, shared, sold, and dispersed around the world. There is a growing consensus across sectors of the need to be aware of, and secure, data. This is more challenging and complex than simple encryption, because data can be aggregated and used to reveal protected or hidden information. Society is recognizing the need for *data privacy* (protection of personal information) and *data security* (protection from malicious actors). Authenticity of data is becoming more difficult to establish, with the advent of disinformation produced by machine learning. The Internet is global and must ensure compliance with new regulations enacted regarding data privacy and accountability. Data privacy, trust, and ethics concerns are heightened due to themes previously explored, including advanced computing, AI, the edge, and IoT.

It is worth differentiating privacy from security. Privacy focuses on the use of personal data, while security focuses on protecting data from malicious attacks and theft. Security has been incorporated as a key concern in system design for at least a decade. While privacy has been considered at times, the growing magnitude and applications of data—particularly about individuals—have swiftly raised its importance.

Subfields of Interest

Differential privacy addresses the challenge of publicly sharing data set information about patterns of groups while withholding individual information. It is important for the census, medical analyses, and other data analysis efforts that involve gathering information about individuals. Differential privacy adds noise to the data in a very prescribed and mathematically rigorous way that preserves the properties of the overall data while hiding individual identities. NIST published a blog post to help enterprises and groups manage differential privacy in July [Near et al. 2020].

Blockchain is a distributed ledger technology with roots in Bitcoin. Blockchain creates pervasive business opportunities by establishing an immutable ledger for recording transactions, tracking assets, building trust, and enabling smart contracts. Hyperledger Fabric, released through the Linux Foundation, has become a leading collaboration mechanism. IBM is making a big push in blockchain as part of its 5in5 strategy. Gartner projects practical enterprise applications in the next three to five years. The US Air Force has some embryonic efforts (funded through the Small Business Innovation Research (SBIR) program) that use Hyperledger Fabric for supply chain logistics.

Trust has many aspects, among which is *confidence* in the data you see or the output of a system—particularly an AI-enabled system. Machine learning is empowering the creation of fraudulent media. This includes deceptive text posted on social media, and synthesized audio and video, commonly referred to as deepfakes. The wide availability of footage and cutting-edge deep learning (particularly Generative Adversarial Networks, or GANs), have made deepfakes substantially easier to produce. Common techniques include face synthesis, identity swap, attribute manipulation, and expression swap. Deepfake detection researchers are embroiled in an arms race with continuously improving deep fake generation technology [Tolosana et al. 2020]. Detection is easier for longer videos (which provide more data) and aided by metadata (e.g., the rest of the social media post or image steganography). The societal challenge is real-time detection of a deepfake—while they can easily be determined after the fact, the disinformation may have already propagated on social media.

Explainable AI concerns the ability to understand why the AI made a given decision. Often, tradeoffs exist between accuracy and explainability. Improving explainability also benefits system qualities such as fairness, testing, and safety. Some key considerations of explainability include:

- Can the AI system explain its answer?
- Can I really have confidence, or is this outcome bias?
- Are there small changes to inputs that would alter the system's predictions?

DARPA launched its Explainable AI (XAI) program in 2017. The program deliverable will be "a toolkit library consisting of machine learning and human-computer interface software modules that could be used to develop future explainable AI systems" [Turek 2017].

Fairness in AI considers whether AI systems produce unfair outcomes. This could result in harms of allocation (e.g., resources or services withheld from certain groups) and harms of representation (e.g., propagation of negative stereotypes). The topic is relevant to contemporary social justice issues, and has attracted much new research. Cathy O'Neil's book *Weapons of Math Destruction* summarizes many of the fairness issues.

Interests

The Confidential Computing Consortium (CCC) is a community focused on securing data using hardware-based trusted execution environment (TEE) technologies and standards. This consortium was established through the Linux Foundation in September, 2019 and includes major technology companies such as Microsoft, Google, Baidu, and Tencent. Their mission is to "address data in use, enabling encrypted data to be processed in memory without exposing it to the rest of the system, reducing exposure to sensitive data and providing greater control and transparency for users" [Linux Foundation 2019].

Commenting on deepfakes, the recent JAIC director, Lt. Gen. Jack Shanahan, said their risk extended beyond civil society: "It's a national security problem as well. We have to invest a lot in it." [Strout 2019]. In a 2020 report, the Integration Division Chief at the US Special Operations Command (SOCOM) declared that violent extremist organizations would use deepfakes "as a cognitive attack vector to sow doubt and confusion among their adversaries to advance their interests and narratives" [Bazin 2020].

Conclusion

This analysis of emerging technologies shows that there is an amazing future ahead of us for innovation, national security, and economic well-being. This future critically depends upon the development of the required software. Advances in software engineering will be needed to support this software production and implementation.

Hardware advances in computing support a wide range of diverse computational architectures. These architectures will require new algorithms to support their implementation, especially at the edge. Power will be a paramount concern. Power-aware computing, a topic of current interest, will become increasingly important. Algorithm developers will have to be extremely conscious of the power requirements to run the software they develop.

A central question is how will this software be developed—will there be sufficient talent? There are hopes for low code/no code scenarios that will enable more people to be able to implement code. However, these low code/no code environments require extensive software engineering to function. AI techniques from the IT giants for developing code are beginning to show some promise, and it is important to understand their potential scope and scale. High-level languages have the potential to increase the effectiveness of advanced programmers.

Data is becoming increasing available, important, and valued. Appropriately using this data while simultaneously protecting it and preventing its misuse presents serious architectural and software engineering challenges.

Disruptions loom on the horizon. Quantum computing is currently capturing national and international attention as one such disruption. The world is currently focused on this Noisy Intermediate Scale Quantum (NISQ) era, where error correction is not capable of being employed. In the near term, NISQ needs to demonstrate success to continue earning industrial support for further development.

References

URLs are valid as of the publication date of this document.

[Amershi et al. 2019]
Amershi, S.; Weld, D.; Vorvoreanu, M.; Fourney, A.; Nushi, B.; Collisson, P.; Suh, J.; Iqbal, S. T.; Bennett, P.; Inkpen, K.; Teevan, J.; Kikin-Gil, R.; and Horvitz, E. Guidelines for Human-AI Interaction. Pages 1-13. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI 19)*. Glasgow, Scotland, UK. May, 2019. https://www.microsoft.com/en-us/research/uploads/prod/2019/01/Guidelines-for-Human-AI-Interaction-camera-ready.pdf

[Army 2020]
U.S. Army CCDC Army Research Laboratory Public Affairs. Army researchers augment combat vehicles with AI. *U.S. Army website*. June 25, 2020. https://www.army.mil/article/236733/army_researchers_augment_combat_vehicles_with_ai

[Bazin 2020]
Bazin, Aaron and Erfourth, Montgomery. Extremism on the Horizon: The Challenges of VEO Innovation and its Impact on Special Operations. *U.S. Army Mad Scientist*. March 19, 2020. https://community.apan.org/wg/tradoc-g2/mad-scientist/m/articles-of-interest/315470

[Black 2020]
Black, Doug. Microsoft's Massive AI Supercomputer on Azure: 285k CPU Cores, 10k GPUs. *HPC Wire website*. May 20, 2020. https://www.hpcwire.com/2020/05/20/microsofts-ai-supercomputer-on-azure-combinations-of-perceptual-domains/

[Carleton et al. 2020]
Carleton, A. D.; Harper, E.; Lyu, M. R.; Eldh, S.; Xie, T.; and Menzies, T. Expert Perspectives on AI. *IEEE Software*. Volume 37. Number 4. July-August, 2020. Pages 87-94. https://ieeexplore.ieee.org/document/9121622

[China 2017]
Department of International Cooperation Ministry of Science and Technology (MOST), P.R. China. New Generation Artificial Intelligence Development Plan. *China Science & Technology Newsletter*. Number 17. September 15, 2017. http://fi.china-embassy.org/eng/kxjs/P020171025789108009001.pdf

[Daugherty et al. 2020]
Daugherty, Paul; Carrel-Billiard, Marc; and Biltz, Michael. Technology Vision 2020. *Accenture website*. December 11, 2020 [accessed]. https://www.accenture.com/us-en/insights/technology/technology-trends-2020

[Dean 2020]
Dean, Jeff. Bringing the benefits of AI to everyone. *Google AI website*. November 24, 2020 [accessed]. https://ai.google/about/

[Deloitte 2020]
Deloitte. Deloitte Insights: Tech Trends 2020. *Deloitte website.* December 11, 2020 [accessed]. https://www2.deloitte.com/content/dam/Deloitte/pt/Documents/tech-trends/TechTrends2020.pdf

[DoD 2018]
Office of the Deputy Assistant Secretary of Defense for Systems Engineering. Department of Defense Digital Engineering Strategy. *Securearc website.* June, 2018. https://sercuarc.org/wp-content/uploads/2018/06/Digital-Engineering-Strategy_Approved.pdf

[Forbes 2020]
Forbes, Stephen. Blackjack. *DARPA website*. November 24, 2020 [accessed]. https://www.darpa.mil/program/blackjack

[Freund 2019]
Freund, Karl. 2019: A Cambrian Explosion In Deep Learning, Part 1. *Forbes website*. January 23, 2019. https://www.forbes.com/sites/moorinsights/2019/01/23/2019-a-cambrian-explosion-in-deep-learning-part-1/

[Gentry 2009]
Gentry, C. and Boneh, D. [Doctoral Diss.] *A fully homomorphic encryption scheme*. Stanford: Stanford university. 2009. https://crypto.stanford.edu/craig/craig-thesis.pdf

[Gidney 2019]
Gidney, C. and Ekerå, M., 2019. How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits. *arXiv*: *Quantum Physics*. 2019. https://arxiv.org/abs/1905.09749

[Gil 2019]
Gil, Y. and Selman, B., 2019. A 20-Year Community Roadmap for Artificial Intelligence Research in the US. *arXiv: Computers and Society*. August 7, 2019. https://arxiv.org/abs/1908.02624

[IBM 2019]
IBM Research. #twinning: Farming's digital doubles will help feed a growing population using less resources. *IBM Research*. November 24, 2020 [accessed]. https://web.archive.org/web/20200206180808/http://www.research.ibm.com/5-in-5/seed/

[Kumar 2020]
Kumar, R.S.S.; Nyström, M.; Lambert, J.; Marshall, A.; Goertzel, M.; Comissoneru, A.; Swann, M.; and Xia, S. Adversarial Machine Learning—Industry Perspectives. *arXiv: Computers and Society*. May 21, 2020. https://arxiv.org/abs/2002.05646

[Linux Foundation 2019]
The Linux Foundation. Confidential Computing Consortium Establishes Formation with Founding Members and Open Governance Structure. *The Linux Foundation website*. October 17, 2019. https://www.linuxfoundation.org/uncategorized/2019/10/confidential-computing-consortium-establishes-formation-with-founding-members-and-open-governance-structure/

[Lu 2020]
Lu, S.; Duan, L.M.; and Deng, D.L. 2020. Quantum adversarial machine learning. *Physical Review Research*. Volume 2. Number 3. August, 2020. https://journals.aps.org/prresearch/abstract/10.1103/PhysRevResearch.2.033212

[Moore 2020]
Moore, Samuel. Cerebras's Giant Chip Will Smash Deep Learning's Speed Barrier. *IEEE Spectrum*. January 1, 2020. https://spectrum.ieee.org/semiconductors/processors/cerebrass-giant-chip-will-smash-deep-learnings-speed-barrier

[Near et al. 2020]
Near, Joseph; Darais, David; and Boeckl, Kaitlin. Differential Privacy for Privacy-Preserving Data Analysis: An Introduction to our Blog Series. *NIST website*. July 27, 2020. https://www.nist.gov/blogs/cybersecurity-insights/differential-privacy-privacy-preserving-data-analysis-introduction-our

[O'Meara 2019]
O'Meara, Sarah. Will China lead the world in AI by 2030? *Nature*. Volume 572. August 21, 2019. Pages 427-428. https://www.nature.com/articles/d41586-019-02360-7

[Reding 2020]
Reding, D.F. and Eaton, J. Science & Technology Trends 2020-2040: Exploring the S&T Edge. *NATO Science & Technology Organization website*. March, 2020. https://www.nato.int/nato_static_fl2014/assets/pdf/2020/4/pdf/190422-ST_Tech_Trends_Report_2020-2040.pdf

[Sculley et al. 2015]
Sculley, David; Holt, Gary; Golovin, Daniel; Davydov, Eugene; Phillips, Todd; Ebner, Dietmar; Chaudhary, Vinay; Young, Michael; Crespo, Jean-Francois; and Dennison, Dan. Hidden technical debt in machine learning systems. Pages 2503–2511. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2 (NIPS'15).* Montreal, Canada. December, 2015. MIT Press, Cambridge, MA, USA. https://proceedings.neurips.cc/paper/2015/file/86df7dcfd896fcaf2674f757a2463eba-Paper.pdf

[Singapore 2020]
Singapore. Virtual Singapore. *National Research Foundation, Prime Minister's Office, Singapore website*. November 24, 2020 [accessed]. https://www.nrf.gov.sg/programmes/virtual-singapore

[Team Orlando 2020]
Team Orlando. Team Orlando Capabilities. *Team Orlando website*. November 24, 2020 [accessed]. http://www.teamorlando.org/team-orlando-capabilities/

[Tolosana 2020]
Tolosana, R.; Vera-Rodriguez, R.; Fierrez, J.; Morales, A.; and Ortega-Garcia, J. Deepfakes and beyond: A survey of face manipulation and fake detection. *arXiv: Computer Vision and Pattern Recognition.* June 18, 2020. https://arxiv.org/abs/2001.00179

[Tung 2020]
Tung, Liam. SpaceX: We've launched 32,000 Linux computers into space for Starlink internet. *ZDNet website*. June 8, 2020. https://www.zdnet.com/article/spacex-weve-launched-32000-linux-computers-into-space-for-starlink-internet/

[Turek 2017]
Turek, Matt. Explainable Artificial Intelligence (XAI). *DARPA website*. November 24, 2020 [accessed]. https://www.darpa.mil/program/explainable-artificial-intelligence

[Vincent 2020]
Vincent, James. OpenAI's latest breakthrough is astonishingly powerful, but still fighting its flaws. *The Verge website*. July 30, 2020. https://www.theverge.com/21346343/gpt-3-explainer-openai-examples-errors-agi-potential

[Weiser 1991]
Weiser, M., The Computer for the 21st Century. *Scientific American*. Volume 265. Number 3. September 1991. Pages 94-105. Reprinted with permission from Scientific American by ACM SIGMOBILE Mobile Computing and Communications Review. July, 1999. https://dl.acm.org/doi/10.1145/329124.329126

Bibliography

URLs are valid as of the publication date of this document.

Emerging Technology Lists

[CCC 2020]
Computing Community Consortium. Technical Focus Areas. *Computing Community Consortium Catalyst website.* November, 20, 2020 [accessed]. https://cra.org/ccc/resources/ccc-by-cs-area/

[Daugherty et al. 2020]
Daugherty, Paul; Carrel-Billiard, Marc; and Biltz, Michael. *Technology Vision 2020*. Accenture. 2020. https://www.accenture.com/us-en/insights/technology/technology-trends-2020

[Deloitte 2020]
Deloitte. Deloitte Insights: Tech Trends 2020. *Deloitte website.* December 11, 2020 [accessed]. https://www2.deloitte.com/content/dam/Deloitte/pt/Documents/tech-trends/TechTrends2020.pdf

[DoD 2019]
Department of Defense. DoD Digital Modernization Strategy: DoD Information Resource Management Strategic Plan FY19-23. *Department of Defense website*. July 5, 2019. https://media.defense.gov/2019/Jul/12/2002156622/-1/-1/1/DOD-DIGITAL-MODERNIZATION-STRATEGY-2019.PDF

[Fretty 2019]
Fretty, Peter. Top 10 Technologies to Watch in 2020. Industry Week *website*. December 4, 2019. https://www.industryweek.com/technology-and-iiot/media-gallery/22028670/top-10-technologies-to-watch-in-2020/slideshow

[Gartner 2020]
Gartner. Top 10 Strategic Technology Trends for 2020. *Gartner Newsroom website*. October 21, 2019. https://www.gartner.com/en/newsroom/press-releases/2019-10-21-gartner-identifies-the-top-10-strategic-technology-trends-for-2020

[IDC 2020]
IDC. IDC FutureScape: Worldwide IT Industry 2020 Predictions. *IDC website*. October, 2019. https://www.idc.com/getdoc.jsp?containerId=US45599219

[IEEE 2019]
IEEE Computer Society. Top 12 Technology Trends for 2020. *IEEE Computer website*. December 10, 2019. https://www.computer.org/press-room/2019-news/ieee-computer-societys-top-12-technology-trends-for-2020

[Marr 2019]
Marr, Bernard. The 7 Biggest Technology Trends In 2020 Everyone Must Get Ready For Now. Forbes *website*. September 30, 2019. https://www.forbes.com/sites/bernardmarr/2019/09/30/the-7-biggest-technology-trends-in-2020-everyone-must-get-ready-for-now/?sh=4f2de6702261

[MIT 2020]
MIT. 10 Breakthrough Technologies 2020. MIT Technology Review. February 26, 2020. https://www.technologyreview.com/10-breakthrough-technologies/2020/

[NITRD 2020]
Networking Information Technology Research and Development Program. FY2021 NITRD Program Component Areas. *NITRD website*. November 30, 2020 [accessed]. https://www.nitrd.gov/subcommittee/NITRD-PCAs-2021.aspx

[Reding 2020]
Reding, D.F. and Eaton, J. Science & Technology Trends 2020-2040: Exploring the S&T Edge. *NATO Science & Technology Organization website*. March, 2020. https://www.nato.int/nato_static_fl2014/assets/pdf/2020/4/pdf/190422-ST_Tech_Trends_Report_2020-2040.pdf

[WEF 2019]
World Economic Forum. Top 10 Emerging Technologies 2019. *World Economic Forum*. June, 2019. http://www3.weforum.org/docs/WEF_Top_10_Emerging_Technologies_2019_Report.pdf

[Y Combinator 2020]
Y Combinator. Requests for Startups. *Y Combinator website*. April, 2020. https://www.ycombinator.com/rfs/

Advanced Computing

[Black 2020]
Black, Doug. Microsoft's Massive AI Supercomputer on Azure: 285k CPU Cores, 10k GPUs. *HPC Wire website*. May 20, 2020. https://www.hpcwire.com/2020/05/20/microsofts-ai-supercomputer-on-azure-combinations-of-perceptual-domains/

[Forbes 2020]
Forbes, Stephen. Blackjack. *DARPA website*. November 24, 2020 [accessed]. https://www.darpa.mil/program/blackjack

[Freund 2019]
Freund, Karl. 2019: A Cambrian Explosion In Deep Learning, Part 1. *Forbes website*. January 23, 2019. https://www.forbes.com/sites/moorinsights/2019/01/23/2019-a-cambrian-explosion-in-deep-learning-part-1/

[Gentry 2009]
Gentry, C. and Boneh, D. [Doctoral Diss.] *A fully homomorphic encryption scheme*. Stanford: Stanford University. 2009. https://crypto.stanford.edu/craig/craig-thesis.pdf

[Gidney 2019]
Gidney, C. and Ekerå, M. How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits. *arXiv*: *Quantum Physics*. 2019. https://arxiv.org/abs/1905.09749

[Goldwasser 2020]
Goldwasser, Shafi. *Cryptography & Machine Learning: What Else?* Presented at Crypto 2020. August, 2020 [virtual conference]. https://crypto.iacr.org/2018/slides/goldwasser_iacr_distinguished_lecture.pdf.

[Hennessy 2019]
Hennessy, John L. and Patterson, David A. A New Golden Age for Computer Architecture. *Communications of the ACM.* Volume 62. Number 2. February, 2019. Pages 48-60. https://cacm.acm.org/magazines/2019/2/234352-a-new-golden-age-for-computer-architecture/fulltext

[Hooker 2020]
Hooker, Sarah, 2020. The hardware lottery (preprint arXiv:2009.06489). *arXiv: Computers and Society*. September 21, 2020. https://arxiv.org/abs/2009.06489

[Martonosi 2018]
Martonosi, Margaret and Roettele, Martin. Next Steps in Quantum Computing: Computer Science's Role. *Computing Community Consortium Catalyst website.* November, 2018. https://cra.org/ccc/wp-content/uploads/sites/2/2018/11/Next-Steps-in-Quantum-Computing.pdf.

[Moore 2020]
Moore, Samuel. Cerebras's Giant Chip Will Smash Deep Learning's Speed Barrier. *IEEE Spectrum website*. January 1, 2020. https://spectrum.ieee.org/semiconductors/processors/cerebrass-giant-chip-will-smash-deep-learnings-speed-barrier

[NASEM 2019]
National Academies of Sciences, Engineering, and Medicine. *Quantum Computing: Progress and Prospects*. The National Academies Press. 2019. ISBN 978-0-309-47969-1.

[NIST 2017]
NIST. Post-Quantum Cryptography. *NIST Computer Security Resource Center website*. November 30, 2020 [accessed]. https://csrc.nist.gov/projects/post-quantum-cryptography

[NIST 2019]
NIST. 26 Algorithms Advancing to the Post-Quantum Crypto 'Semifinals.' *NIST website*. January 30, 2019. https://www.nist.gov/news-events/news/2019/01/nist-reveals-26-algorithms-advancing-post-quantum-crypto-semifinals

[NSTC 2018]
National Science and Technology Council, Committee on Science, Subcommittee on Quantum Information Science. National Strategic Overview for Quantum Information Science. *Executive Office of the President of the United States website*. September, 2018. https://www.whitehouse.gov/wp-content/uploads/2018/09/National-Strategic-Overview-for-Quantum-Information-Science.pdf

The Smarter Edge

[Campbell 2019]

Campbell, Mark. Smart Edge: The Effects of Shifting the Center of Data Gravity Out of the Cloud. *IEEE Computer.* Volume 52. Number 12. December, 2019. Pages 99-102. https://doi.ieeecomputersociety.org/10.1109/MC.2019.2948248

[Roy 2019]

Roy, K.; Jaiswal, A.; and Panda, P. Towards spike-based machine intelligence with neuromorphic computing. *Nature.* Volume 575. November, 2019. Pages 607–617. https://doi.org/10.1038/s41586-019-1677-2

[Satyanarayanan 2017]

Satyanarayanan, Mahadev. The Emergence of Edge Computing. *IEEE Computer*. Volume 50. Number 1. January, 2017. Pages 30-39. doi: 10.1109/MC.2017.9. https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=7807196

[Tung 2020]

Tung, Liam. SpaceX: We've launched 32,000 Linux computers into space for Starlink internet. *ZDNet website*. June 8, 2020. https://www.zdnet.com/article/spacex-weve-launched-32000-linux-computers-into-space-for-starlink-internet/

[Weiser 1991]

Weiser, M., The Computer for the 21st Century. *Scientific American*. Volume 265. Number 3. September 1991. Pages 94-105. Reprinted with permission from Scientific American by ACM SIGMOBILE Mobile Computing and Communications Review. July, 1999. https://dl.acm.org/doi/10.1145/329124.329126

Digital Twins

[Alber, 2019]

Alber, M.; Buganza Tepole, A.; Cannon, W.R. et al. Integrating machine learning and multiscale modeling—perspectives, challenges, and opportunities in the biological, biomedical, and behavioral sciences. *npj Digital Medicine website*. Volume 2. Number 115. November, 2019. https://doi.org/10.1038/s41746-019-0193-y

[Dambrot 2018]

S. M. Dambrot; de Kerchove, D.; Flammini, F.; Kinsner, W. L.; MacDonald, G.; and Saracco, R. Symbiotic Autonomous Systems: An FDC initiative. White Paper II. *IEEE Symbiotic Autonomous Systems website*. October, 2018. https://digitalreality.ieee.org/images/files/pdf/SAS-WP-II-2018-Finalv3.2.pdf

[DoD 2018]

Department of Defense. Digital Engineering Strategy. *Office of the Deputy Assistant Secretary of Defense for Systems Engineering website*. June, 2018. https://sercuarc.org/wp-content/uploads/2018/06/Digital-Engineering-Strategy_Approved.pdf

[IBM 2019]
IBM Research. #twinning: Farming's digital doubles will help feed a growing population using less resources. *IBM*. November 24, 2020 [accessed]. https://web.archive.org/web/20200206180808/http://www.research.ibm.com/5-in-5/seed/
[*IBM has another list for the new year, so this is an archive from 2019*.]

[Leopold 2020]
Leopold, George. Military Enlists Digital Twin Technology to Secure Chips. *EE Times website*. January 2, 2020. https://www.eetimes.com/military-enlists-digital-twin-technology-to-secure-chips/?utm_source=newsletter&utm_campaign=link&utm_medium=EETimesMilAero-20200108#

[Mussomeli 2020]
Mussomeli, Adam; Parrott, Aaron; Umbenhauer, Brian; and Warshaw, Lane. Digital twins: Bridging the physical and digital. *Deloitte Insights website*. January 15, 2020. https://www2.deloitte.com/us/en/insights/focus/tech-trends/2020/digital-twin-applications-bridging-the-physical-and-digital.html

[Purdy 2020]
Purdy, Mark; Eitel-Porter, Ray; Krüger, Robert; and Deblaere, Thijs. "How Digital Twins Are Reinventing Innovation." MIT Sloan Management Review. January 14, 2020. https://sloanreview.mit.edu/article/how-digital-twins-are-reinventing-innovation/

[Saracco 2019a]
R. Saracco. Digital Twins: Bridging Physical Space and Cyberspace. *IEEE Computer*. Volume 52. Number 12. December, 2019. Pages 58-64. doi: 10.1109/MC.2019.2942803 https://ieeexplore.ieee.org/abstract/document/8909940

[Saracco 2019b]
R. Saracco. Digital Twins: Where We Are Where We Go – II. *IEEE Future Directions website*. November 30, 2020 [accessed]. https://cmte.ieee.org/futuredirections/2019/06/30/digital-twins-where-we-are-where-we-go-ii/

[Singapore 2020]
Singapore. Virtual Singapore. *National Research Foundation, Prime Minister's Office, Singapore website*. November 24, 2020 [accessed]. https://www.nrf.gov.sg/programmes/virtual-singapore

Artificial Intelligence

[Amazon 2020]
Amazon. Amazon CodeGuru. *Amazon Web Services website*. August 20, 2020 [accessed]. https://aws.amazon.com/codeguru/

[Amershi 2019]
Amershi, Saleema et al. Software Engineering for Machine Learning: A Case Study. Pages 291-300. In *Proceedings of the 41st International Conference on Software Engineering: Software Engineering in Practice*. Montreal, Quebec, Canada. May, 2020. https://doi.org/10.1109/ICSE-SEIP.2019.00042

[Amershi et al. 2019]
Amershi, S.; Weld, D.; Vorvoreanu, M.; Fourney, A.; Nushi, B.; Collisson, P.; Suh, J.; Iqbal, S. T.; Bennett, P.; Inkpen, K.; Teevan, J.; Kikin-Gil, R.; and Horvitz, E. Guidelines for Human-AI Interaction. Pages 1-13. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI 19)*. Glasgow, Scotland, UK. May, 2019. https://www.microsoft.com/en-us/research/uploads/prod/2019/01/Guidelines-for-Human-AI-Interaction-camera-ready.pdf

[Army 2020]
U.S. Army CCDC Army Research Laboratory Public Affairs. Army researchers augment combat vehicles with AI. *U.S. Army website*. June 25, 2020. https://www.army.mil/article/236733/army_researchers_augment_combat_vehicles_with_ai

[Bergstein 2020]
Bergstein, Brian. What AI still can't do. *MIT Technology Review website.* February 19, 2020. https://www.technologyreview.com/2020/02/19/868178/what-ai-still-cant-do/

[Breck 2017]
Breck, E.; Cai, S.; Nielsen, E.; Salib M.; and Sculley, D. The ML test score: A rubric for ML production readiness and technical debt reduction. Pages 1123-1132. In *Proceedings of 2017 IEEE International Conference on Big Data (Big Data)*. Boston, MA. December, 2017. doi: 10.1109/BigData.2017.8258038. https://ieeexplore.ieee.org/document/8258038

[Carleton et al. 2020a]
Carleton, A. D.; Harper, E.; Menzies, T.; Xie, T.; Eldh S.; and Lyu, M. R. The AI Effect: Working at the Intersection of AI and SE. *IEEE Software*. Volume 37. Number 4. July-August, 2020. Pages 26-35. doi: 10.1109/MS.2020.2987666. https://ieeexplore.ieee.org/document/9121618

[Carleton et al. 2020b]
Carleton, A. D.; Harper, E.; Lyu, M. R.; Eldh, S.; Xie, T.; and Menzies, T. Expert Perspectives on AI. *IEEE Software*. Volume 37. Number 4. July-August, 2020. Pages 87-94. https://ieeexplore.ieee.org/document/9121622

[China 2017]
Department of International Cooperation Ministry of Science and Technology (MOST), P.R. China. New Generation Artificial Intelligence Development Plan. *China Science & Technology Newsletter*. Number 17. September 15, 2017. http://fi.china-embassy.org/eng/kxjs/P020171025789108009001.pdf

[Dawson 2019]
Dawson, Caitlin. Glimpsing into the Future of AI: A Conversation with Yolanda Gil. *USCViterbi website*. September 18, 2019. https://viterbischool.usc.edu/news/2019/09/glimpsing-into-the-future-of-ai-a-conversation-with-yolanda-gil/

[Dean 2020]
Dean, Jeff. Bringing the benefits of AI to everyone. *Google AI website*. November 24, 2020 [accessed]. https://ai.google/about/

[Gil 2019]
Gil, Y. and Selman, B., 2019. A 20-Year Community Roadmap for Artificial Intelligence Research in the US. *arXiv: Computers and Society*. August 7, 2019. https://arxiv.org/abs/1908.02624

[Kästner 2020a]
Kästner, Christian. Software Engineering for AI-Enabled Systems (CMU 17-445/645, Summer 2020). *GitHub website*. November 30, 2020 [accessed]. https://ckaestne.github.io/seai/S2020/

[Kästner 2020b]
Kästner, Christian. Software Engineering for AI/ML—An Annotated Bibliography. *GitHub website*. August 20, 2020 [accessed]. https://github.com/ckaestne/seaibib

[Kumar 2020]
Kumar, R.S.S.; Nyström, M.; Lambert, J.; Marshall, A.; Goertzel, M.; Comissoneru, A.; Swann, M.; and Xia, S. Adversarial Machine Learning—Industry Perspectives. *arXiv: Computers and Society*. May 21, 2020. https://arxiv.org/abs/2002.05646

[Lu 2020]
Lu, S.; Duan, L.M.; and Deng, D.L. 2020. Quantum adversarial machine learning. *Physical Review Research*. Volume 2. Number 3. August, 2020. https://journals.aps.org/prresearch/abstract/10.1103/PhysRevResearch.2.033212

[Marco Polo 2020]
Marco Polo. The Global AI Talent Tracker. *Marco Polo website*. August 20, 2020 [accessed]. https://macropolo.org/digital-projects/the-global-ai-talent-tracker/?mod=djemDailyShot&mod=djemDailyShot

[O'Meara 2019]
O'Meara, Sarah. Will China lead the world in AI by 2030? *Nature*. Volume 572. August 21, 2019. Pages 427-428. https://www.nature.com/articles/d41586-019-02360-7

[Ozkaya 2020]
Ozkaya, Ipek. What Is Really Different in Engineering AI-Enabled Systems? *IEEE Software*. Volume 37. Number 4. July-August, 2020. Pages 3-6. doi: 10.1109/MS.2020.2993662. https://www.computer.org/csdl/magazine/so/2020/04/09121629/1kMT6KOMvzG

[Sculley et al. 2015]
Sculley, David; Holt, Gary; Golovin, Daniel; Davydov, Eugene; Phillips, Todd; Ebner, Dietmar; Chaudhary, Vinay; Young, Michael; Crespo, Jean-Francois; and Dennison, Dan. Hidden technical debt in machine learning systems. Pages 2503–2511. In *Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2 (NIPS '15)*. Montreal, Canada. December, 2015. MIT Press, Cambridge, MA, USA. http://papers.nips.cc/paper/5656-hidden-technical-debt-in-machine-learning-systems.pdf

[Simonite 2020]
Simonite, Tom. OpenAI's Text Generator Is Going Commercial. *Wired website*. June 11, 2020. https://www.wired.com/story/openai-text-generator-going-commercial/

[Ning 2020]
Ning, Emma. Microsoft open sources breakthrough optimizations for transformer inference on GPU and CPU [blog post]. *Microsoft Open Source Blog*. January 21, 2020. https://cloudblogs.microsoft.com/opensource/2020/01/21/microsoft-onnx-open-source-optimizations-transformer-inference-gpu-cpu/

[Vincent 2020a]
Vincent, Brandi. Legislation launched in early June received new attention and support this week. *NextGov website*. July 2, 2020. https://www.nextgov.com/emerging-tech/2020/07/congress-inches-closer-creating-national-cloud-ai-research/166624/

[Vincent 2020b]
Vincent, James. OpenAI's latest breakthrough is astonishingly powerful, but still fighting its flaws. *The Verge website*. Jul 30, 2020. https://www.theverge.com/21346343/gpt-3-explainer-openai-examples-errors-agi-potential

Extended Reality

[Apple 2019]
Apple, Inc. Apple Offers New Augmented Reality Art Sessions. *Apple website*. July 30, 2019. https://www.apple.com/newsroom/2019/07/apple-offers-new-augmented-reality-art-sessions

[Bayern 2020]
Bayern, Macy. 10 mixed reality trends to expect in 2020. *TechRepublic website*. January 31, 2020. https://www.techrepublic.com/article/10-mixed-reality-trends-to-expect-in-2020/

[BISim 2020]
Bohemia Interactive Simulations (BISim). *VBS4 website.* August 20, 2020 [accessed]. http://www.bisimulation.com

[Evans 2020]
Evans, Jon. Magic Leap's $2.6 billion bait and switch. *TechCrunch website*. April 26, 2020. https://techcrunch.com/2020/04/26/tragic-leap/?guccounter=1

[FIG 2020]
Carnegie Mellon University. Future Interfaces Group. August 20, 2020 [accessed]. http://www.figlab.com/

[Hodicky 2015]
Hodicky, Jan (ed.). *Modelling and Simulation for Autonomous Systems: Second International Workshop (MESAS 2015): Revised Selected Papers*. Prague, Czech Republic. April 29-30, 2015. Springer International Publishing. 2015.

[Lange 2020]
Lange, Katie. Virtual, Augmented Reality Are Moving Warfighting Forward [blog post]. *Inside DoD blog, Department of Defense website*. February 10, 2020. https://www.defense.gov/Explore/Inside-DOD/Blog/Article/2079205/how-virtual-augmented-reality-are-moving-warfighting-forward/

[Marr 2019]
Marr, Bernard. What Is Extended Reality Technology? A Simple Explanation For Anyone. *Forbes website*. August 12, 2019. https://www.forbes.com/sites/bernardmarr/2019/08/12/what-is-extended-reality-technology-a-simple-explanation-for-anyone/#5c1669d37249

[Schmalstieg 2016]
Schmalstieg, Dieter and Höllerer, Tobias. *Augmented Reality: Principles and Practice*. O'Reilly Media, Inc. 2016. ISBN: 9780133153217. https://learning.oreilly.com/library/view/augmented-reality-principles/9780133153217/cover.html

[SCMP 2019]
South China Morning Post. These Star Trek-Style Translation Devices Are Smashing the Language Barrier, Thanks to AI and Cloud Computing. *South China Morning Post*. January 11, 2019. https://www.scmp.com/news/world/united-states-canada/article/2181601/these-near-real-time-translation-devices-are

[Team Orlando 2020]
Team Orlando. Team Orlando Capabilities. *Team Orlando website*. November 24, 2020 [accessed]. http://www.teamorlando.org/team-orlando-capabilities/

[Visbox 2020a]
Visbox, Inc. CAVE Automatic Virtual Environment. *Visbox, Inc. website.* August 20, 2020 [accessed]. http://www.visbox.com/products/cave/

[Visbox 2020b]
Visbox, Inc. Immersive 3D Applications. *Visbox, Inc. website.* August 20, 2020 [accessed]. http://www.visbox.com/applications/immersive-3d/

[Wyoming 2020]
University of Wyoming School of Energy Resources. Shell 3D Virtualization Center. *University of Wyoming website*. August 20, 2020 [accessed]. https://www.uwyo.edu/ser/visualization-center/

Data Privacy, Trust, and Ethics

[Mendelson 2019]
A. Mendelson. Security and Privacy in the Age of Big Data and Machine Learning. *IEEE Computer*. Volume 52. Number 12. Dec. 2019. Pages 65-70. doi: 10.1109/MC.2019.2943137. https://ieeexplore.ieee.org/document/8909933

Ethical Use of (Disruptive) Technology

[Bannister 2020]

Bannister, Catherine and Golden, Deborah. Ethical technology and trust: Applying your company's values to technology, people, and processes. *Deloitte Insights website*. January 15, 2020. https://www2.deloitte.com/us/en/insights/focus/tech-trends/2020/ethical-technology-and-brand-trust.html

[Barber, 2019]

Barber, Gregory and Molteni, Megan. Google Is Slurping Up Health Data—and It Looks Totally Legal. *Wired website*. November 11, 2019. https://www.wired.com/story/google-is-slurping-up-health-dataand-it-looks-totally-legal/

[Murphy 2019]

Murphy, Timothy; Garg, Swati; Sniderman, Brenna; and Buckley, Natasha. Ethical technology use in the Fourth Industrial Revolution CEO leadership needed. *Deloitte Insights website*. July 15, 2019. https://www2.deloitte.com/us/en/insights/focus/industry-4-0/ethical-technology-use-fourth-industrial-revolution.html

Differential Privacy

[Dwork 2011]

Dwork, Cynthia. A firm foundation for private data analysis. *Communications of the ACM* Volume 54. Number 1. 2011. Pages 86-95. https://dl.acm.org/doi/10.1145/1866739.1866758

[Hassan 2020]

Hassan, M. U.; Rehmani, M. H.; and Chen, J. Differential Privacy Techniques for Cyber Physical Systems: A Survey. *IEEE Communications Surveys & Tutorials*. Volume 22. Number 1. First quarter 2020. Pages 746-789. doi: 10.1109/COMST.2019.2944748. https://ieeexplore.ieee.org/document/8854247

[Near et al. 2020]

Near, Joseph; Darais, David; and Boeckl, Kaitlin. Differential Privacy for Privacy-Preserving Data Analysis: An Introduction to our Blog Series. *NIST website*. July 27, 2020. https://www.nist.gov/blogs/cybersecurity-insights/differential-privacy-privacy-preserving-data-analysis-introduction-our

Deepfakes

[Bazin 2020]

Bazin, Aaron and Erfourth, Montgomery. Extremism on the Horizon: The Challenges of VEO Innovation and its Impact on Special Operations. *U.S. Army Mad Scientist*. March 19, 2020. https://community.apan.org/wg/tradoc-g2/mad-scientist/m/articles-of-interest/315470

[Chesney 2019]
Chesney, Bobby and Citron, Danielle. Deep fakes: a looming challenge for privacy, democracy, and national security. *Calif. Law Review.* Volume 107. 2019. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3213954

[Strout 2019]
Strout, Nathan. How the Pentagon is tackling deepfakes as a national security problem. *C4ISRNET website*. August 29, 2019. https://www.c4isrnet.com/information-warfare/2019/08/29/how-the-pentagon-is-tackling-deepfakes-as-a-national-security-problem/

[Tolosana 2020]
Tolosana, R.; Vera-Rodriguez, R.; Fierrez, J.; Morales, A.; and Ortega-Garcia, J. Deepfakes and beyond: A survey of face manipulation and fake detection. *arXiv: Computer Vision and Pattern Recognition.* June 18, 2020. https://arxiv.org/abs/2001.00179

Others

[Wheatly 2020]
Wheatly, Mike. Accenture, AMD, Facebook and Nvidia sign up to advance 'Confidential Computing.' *SiliconANGLE website*. June 29, 2020. https://siliconangle.com/2020/06/29/accenture-amd-facebook-nvidia-sign-advance-confidential-computing/

[Kaplan 2016]
Kaplan, Fred. 'WarGames' and Cybersecurity's Debt to a Hollywood Hack. *New York Times*. February 21, 2016. https://www.nytimes.com/2016/02/21/movies/wargames-and-cybersecuritys-debt-to-a-hollywood-hack.html

[Kozyrkov 2020]
Kozyrkov, Cassie. Focus on decisions, not outcomes! *Towards Data Science* (*Medium) website*. May 17, 2020. https://towardsdatascience.com/focus-on-decisions-not-outcomes-bf6e99cf5e4f

[Linux Foundation 2019]
The Linux Foundation. Confidential Computing Consortium Establishes Formation with Founding Members and Open Governance Structure. *The Linux Foundation website*. October 17, 2019. https://www.linuxfoundation.org/uncategorized/2019/10/confidential-computing-consortium-establishes-formation-with-founding-members-and-open-governance-structure/

[Turek 2017]
Turek, Matt. Explainable Artificial Intelligence (XAI). *DARPA website*. November 24, 2020 [accessed]. https://www.darpa.mil/program/explainable-artificial-intelligence

Contact Us

Software Engineering Institute
4500 Fifth Avenue, Pittsburgh, PA 15213-2612

Phone: 412.268.5800 | 888.201.4479
Web: www.sei.cmu.edu
Email: info@sei.cmu.edu

This material is based upon work funded and supported by the Department of Defense under Contract No. FA8702-15-D-0002 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center.

DM20-1165