Federal Register Notice 88 FR 7999, Federal Register: Request for Information on the 2023 Federal Cybersecurity Research and Development Strategic Plan, March 3rd, 2023

Request for Information on the 2023 Federal Cybersecurity Research and Development Strategic Plan

Computing Research Association (CRA)'s Computing Community Consortium (CCC)

Computing Community Consortium's Response to RFI on the 2023 Federal Cybersecurity Research and Development Strategic Plan

March 3, 2023

Written by: *Nadya Bliss (Arizona State University), Elizabeth Bradley (University of Colorado-Boulder), Randal Burns (Johns Hopkins University), Thomas M. Conte (Georgia Institute of Technology), David Danks (University of California San Diego), Nathan Evans (Arizona State University), Kevin Fu (Northeastern University), Haley Griffin (Computing Community Consortium), William D. Gropp (University of Illinois Urbana-Champaign), David Jensen (University of Massachusetts Amherst), Chandra Krintz (University of California-Santa Barbara), Brian LaMacchia (Farcaster Consulting Group), Daniel Lopresti (Lehigh University), Madhav Marathe, (University of Virginia), Melanie Moses (University of New Mexico), Ann W. Schwartz (Computing Community Consortium), Ufuk Topcu (University of Texas-Austin), and Pamela Wisniewski (Vanderbilt University)*

This response is from the Computing Research Association (CRA)'s Computing Community Consortium (CCC). CRA is an association of nearly 250 North American computing research organizations, both academic and industrial, and partners from the professional societies. The mission of the CCC is to bring together the computing research community to enable the pursuit of innovative, high-impact computing research that aligns with pressing national and global challenges.

The Appendix to this response contains both the reference to the CCC Response to NITRD "RFI on Update to the 2016 Federal Cybersecurity Research and Development Strategic Plan" (2019) and the relevant white papers and visioning activities since the last response.

This response includes answers to questions 1, 3, 4, 5, 6 and 7. Detailed responses are below, but we highlight a few key points here:

- **Socio-technical resilience and human aspects of cyber security:** While the 2019 Federal Strategy did include multiple references to the importance of the Human Aspects of cybersecurity, it did not elevate that topic to a “Priority Area”. While it is, of course, a “Critical Dependency” as it is referred to in the previous strategy, it is also important to invest in R&D in this area. Research in this area could potentially include development of multi-scale multi-theory models to understand interdependent socio-technical infrastructure systems. This can lead to identification of new vulnerabilities, making systems more resilient, early warning systems, and understanding inter-dependencies. A big challenge is the availability of data and including this in the federal cybersecurity strategy can initiate new ways data can be shared safely.
- **Resilience and security by design:** The 2019 Federal Strategy has significant focus on cyber defense. A key theme of the comments below is to incorporate security up front, by design and not as an afterthought. That, together with socio-technical resilience as described above, is likely to lead to more secure systems.
- **Artificial intelligence:** Recently, there has been increased adoption of Large Language Models (LLMs) and generative AI models in general. These potentially present a significant cybersecurity risk, particularly in their ability to generate disinformation effectively and efficiently, and at an overwhelming scale. More broadly, the ability to discern between authentic, accurate, auto-generated, and maliciously generated information via artificial intelligence (regardless of modality - text, images, video, etc.) presents significant challenges to cybersecurity and needs to be prioritized in the updated research strategy. While the technology companies are increasingly investing in dis- and misinformation related work, this work needs to be continually complemented by academic research and education initiatives.
- **Pandemic and computing:** The COVID-19 pandemic led to rapid adoption of remote working environments which continue to persist because both employees and employers find them attractive. Computing capabilities enable significant connectivity and productivity, but also have the potential to lead to a broader attack surface.
- **Climate and computing:** Rapidly accelerating effects of climate change require new research in resiliency and security of computing infrastructure, particularly in context of the accelerating rate of extreme weather events. There are many opportunities for highly impactful computing research in hardware, software, and algorithms that could support both security and efficiency in a co-optimized fashion.

1. What new innovations have the potential to greatly enhance the security, reliability, resiliency, trustworthiness, and privacy protections of the digital ecosystem (including but not limited to data, computing, networks, cyber-physical systems, and participating entities such as people and organizations)?

Computing researchers in the public and private sectors are rapidly innovating to improve digital systems, and there are many instances of great successes that have resulted from these efforts:

Artificial Intelligence/Large Language Models

Continued advances in machine learning have made artificial intelligence a powerful tool in detecting cyber attacks, particularly advanced persistent threats (APTs) that are long lasting, adaptive, and have a small attack signature. Large language models and graphical neural networks can detect anomalies in configurations, programs, scripts, and network traffic. Learned models can be used to characterize normal activity and, as a consequence, detect malicious and faulty data and programs beyond the capabilities of human analysts, rule-based systems, and classical statistical methods. Finally, LLMs and other AI technologies can be deployed to redirect and distract human attackers by providing believable decoy activity, documents, etc; for example the IARPA (Intelligence Advanced Research Projects Activity) program cited below[1]. LLMs and other AI technologies could substantially enhance this approach to cyber-defense.

It is worthwhile to point out that while AI has the potential to enhance cybersecurity, it can also be deployed as a threat accelerator. For example, large language models could be leveraged as either instruments of disinformation or efficiently probing socio-technical systems for weaknesses at scale.

Human-centered privacy frameworks

Examples include Privacy as Contextual Integrity[2], Privacy by Design[3], and Safety by Design[4]. These frameworks take more nuanced/contextualized approaches to designing digital systems in a proactive way that acknowledges that humans are often the biggest threat to security, reliability, resiliency, trustworthiness, and privacy protections.

---

[1] https://www.iarpa.gov/research-programs/rescind
[2]https://digitalcommons.law.uw.edu/wlr/vol79/iss1/10/#:~:text=Contextual%20integrity%20ties%20adequate%20protection,norms%20of%20distribution%20within%20it.
[3] https://link.springer.com/chapter/10.1007/978-3-030-82786-1_2
[4] https://www.datocms-assets.com/22233/1652864615-child-safety-by-design-report-final-1.pdf

Post-deployment code repair

Research in automated and human assisted code repair on legacy codes has also presented opportunities to mitigate cybersecurity risk in deployed systems, thus allowing increased trustworthiness of infrastructure.

Privacy enhancing technologies

Advances in privacy enhancing technologies and their adoption from multi-party computation to homomorphic encryption to privacy preserving video analytics (as for example, in Homeland Security relevant scenarios[5]) has also been observed in the last few years.

3. What areas of research or topics of the 2019 Strategic Plan should continue to be a priority for federally funded research and require continued Federal R&D investments?

Re-framing: Human-Centered Design, Resilience and Empowerment

The areas of research in the 2019 Strategic Plan are largely missing a critical part of cybersecurity: humans. In this context, people are both part of the solution, the problem, and an entity that needs to be secured. It is critical that continued Federal R&D investments have a sociocultural / sociotechnical lens that accommodates the people using and impacted by digital systems.

Creating resilient systems requires creative solutions that consider the humans interacting within the system. One way the topics could be adjusted to better accommodate humans is to include “Human-centered Design” or “Using Human-centeredness to Design Complex Socio-Technical Systems” as a cross-cutting priority. Human-centeredness is necessary for AI (e.g., fairness, bias, explainability) and the other more technical approaches to be effective. Human behavior is often a weak point in cybersecurity systems, for example when authorized users fall prey to phishing and other attacks. The creation of resilient systems also needs to take into account economic realities and other implementation or adoption challenges. If funding is not

[5] https://pets4hse.org/index.html

available to widely implement a system that already exists, like underfunded public hospitals lacking the resources to protect patient records from cyber criminals, then that system is not an effective or resilient solution and further research is necessary.

Broadly, the 2019 Strategic Plan also has a stark focus on defensive elements (i.e., deter, protect, detect, respond), that could be improved by taking more resilience-based or empowering approaches that allow for struggle, failure, and recovery. And privacy is sometimes interpreted too narrowly in the document, while it should be spread throughout the ecosystem and lifetime of systems.

Quantum Information Science

Quantum information science remains a key priority area of investment because of its potential to break cryptosystems. The technology today is most advanced with superconducting quantum technology, but ion-trap and neutral atom technologies are advancing rapidly, though none of these are yet on the type of exponential growth scale that enabled advances in CMOS semiconductors. At the same time, there is a perception that quantum computing has over-promised and under-delivered. This could result in a “quantum winter” much like the “AI winter” that prevented the US from realizing today’s AI technologies years earlier. It is imperative to remain focused on and continue to invest in quantum information science in general and quantum technology for computing specifically through this potential period of skepticism.

One issue that remains understudied in quantum computing is the need for higher level ways to program these systems. Current approaches, even those that use high level programming languages, express computation at a low quantum circuit level. New algorithms and applications for quantum computing are reserved for the domain of theoretical mathematicians and physicists. An analogy is to VLSI technology before the design rule revolution brought by Carver Mead and Lynn Conway: VLSI circuits were designed by applied physicists solving complex electrical equations. Mead and Conway introduced a series of design rules enabling non-specialists to readily create complex VLSI circuits, and as a result, the applications of integrated circuits exploded. A similar renaissance is required in the quantum domain to enable advances and innovation by a larger number of individuals.

4. What objectives not included in the 2019 Strategic Plan should be strategic priorities for federally funded R&D in cybersecurity? Discuss the challenges, desired capabilities and outcomes, and objectives that should guide research to achieve the desired capabilities, and why those capabilities and outcomes should be strategic priorities for federally funded R&D.

Operational Technology Cybersecurity Research

One gap in the 2019 Strategic Plan is Operational Technology (OT) Cybersecurity Research. The key challenge is how to co-design OT hardware and software to remain highly available and integrity protected despite constantly shifting threats that require updates and patches to both software and hardware. The R&D gap is so great that the industry has created the term "legacy device" to describe OT systems that are not only insecure because of known vulnerabilities, but also impossible to secure because of the lack of a mechanism to patch security holes. A desired capability is to enable security updates for legacy OT systems as well as finding economic models that incentivize the replacement of legacy OT systems with systems that can quickly adapt to shifting threat models. A specific outcome could be hospitals no longer needing to turn away patients when ransomware breaks into an OT system. A more general outcome is higher consumer confidence that OT systems will continue to provide essential functions even when under attack. This requires a shift in thinking away from component-centric security modules to system-centric trustworthy systems. These R&D priorities are important not only for science and engineering, but also for society that depends on highly available emergency medicine, reliable power grids, safe transportation, and every day public infrastructure despite endemic cybersecurity threats to availability and integrity of OT systems and their sensors.

Human-centered design and evaluation

Human-centered design and evaluation of complex socio-technical systems should not be conflated with privacy. They should be two separate priorities with the privacy section focused primarily on information flows. The human-centered priority should be cross-cutting (maybe a social, behavioral, and economic sciences orientation) to show that our federal government acknowledges that all of the technical solutions have significant impact on humans.

For solutions to be practical, responsible, and feasible, human-centered design principles must be implemented across national priorities that invoke translational and implementation science. Simply put, sometimes high-tech solutions are not the answer. True translational and implementation science needs to be led by those in the field (e.g.,

non-profits, practitioners, end users), so that the technologies we build to solve these problems do not have unintended consequences and/or negative impacts[6].

Core challenges to designing complex and resilient socio-technical systems include:

- Formative assessments that identify key human problems that need to be addressed, rather than prescribing state-of-the-art technologies to solve these problems without first clearly understanding them.
- Further, summative assessments of how these technologies affect humans once they are deployed are also necessary.
- Team dynamics of interdisciplinary research. Interdisciplinary teams often develop stratification, with subject matter experts in social sciences or humanities often playing a secondary role to the technical team. New incentives to create a more equal balance or even place the technical solutions architects in the 'service role' with the human subject matter experts as the 'customer'.

While such cross-disciplinary approaches are less straightforward and harder to implement, they are what is necessary to affect real positive change. A possible way to do this is to have grant funding mechanisms that are phased in two parts - where the social/behavioral team first scopes the problem space and defines the requirements (agnostic of a particular technical solution). Then, this team could solicit phase 2 proposals from technical teams that they feel would meet these requirements. A third phase would be the implementation/translational science piece, where the solutions that are built are evaluated for feasibility and impact. This approach mirrors more closely what is done in industry, where the customer/subject matter expert (SME) drives the requirements and the development team acts to provide the solution. In academia and when creating grant-funding mechanisms, this hierarchy is often flipped, where the technical experts take lead.

Privacy

While privacy is mentioned in the 2019 Strategic Plan, it is defined too narrowly, with a focus on individual privacy. The Strategic Plan should seek to address privacy problems beyond the individual, including interpersonal privacy threats such as sextortion (e.g., unauthorized sharing of digital imagery of a person) or the distribution of child pornography (e.g., digital rights of youth[7]).

[6] https://arxiv.org/ftp/arxiv/papers/2112/2112.09544.pdf
[7] https://journals.sagepub.com/doi/abs/10.1177/1461444816686318

Another area that could benefit from federal R&D investment is at the intersection of privacy preserving/enhancing technologies and incentives and policies to adopt those technologies. The U.S. has historically been misaligned/behind on regulations and policies that have been designed to protect individual privacy (e.g., not being part of GDPR, COPPA regulations for children being out-of-date). While creating policy would be outside the scope of federally funded research, research on adoption and incentive pathways would not be.

These are two privacy challenges that do not seem to be addressed in the strategic plan. And, while several privacy frameworks already exist, regulations and policies that hold platforms accountable for implementing these frameworks in a consistent way are lagging behind. Research initiatives that support validating which frameworks are useful in what contexts and finding ways to translate theoretical frameworks into actionable design principles and practices have potential to have significant impact.

Security and trustworthiness of the AI ecosystem

While AI is also mentioned in the 2019 Strategic Plan, the gap has widened between how AI-enabled functionality is developed and deployed and what a *trustworthy* distributed digital infrastructure can support today and possibly ever. The vulnerabilities due to potential deployment of large models over a distributed infrastructure are unclear. The training and execution and interactions with these models expand the attack surfaces of the systems in which they are going to be integrated to unprecedented levels. Additionally, the innovation ecosystem has grown to favor empirical performance over principles and guidelines for trustworthiness that have been established over decades-long experience. Deferring the security concerns to late stages of development has never been sustainable, and the exponentially growing complexity and cost—in data, computing, and engineering—amplifies the importance of security in the early stages of development and even conception of technology. It is therefore critical to incentivize security in the entire AI ecosystem. Otherwise, it will be increasingly costly and possibly catastrophic to walk the thin line that separates AI being an opportunity for security from AI being a security vulnerability.

Cryptographic agility and the transition to quantum-resistant cryptography

Finally, while the 2019 Strategic Plan briefly mentions goals of drafting standards for quantum-resistant (a.k.a. post-quantum) cryptography and implementing them, it is silent on all the other components necessary to effect a wide-scale cryptographic

algorithm transition to post-quantum cryptography. As we have seen over the last 18 years (since NSA announced “Suite B” and their goal of transitioning public-key cryptography from RSA to elliptic curve), cryptographic algorithm transitions are difficult, time-consuming, and take much longer than expected. Indeed, many sectors of industry have still not transitioned to elliptic curve cryptography and are now being told to start another transition. The 2023 Strategic Plan should include R&D goals concerning integrating quantum-resistant algorithms into every security protocol that utilizes cryptography and improving system architectures with cryptographic agility to make this PQC transition (and future cryptographic algorithm transitions) easier[8].

5. What other scientific, technological, economic, legal, or societal changes and developments occurring now or in the foreseeable future have the potential to significantly disrupt our abilities to secure the digital ecosystem and make it resilient? Discuss what federally funded R&D could improve the understanding of such developments and improve the capabilities needed to mitigate against such disruptions.

LLMs (and other generative models) vulnerabilities to cyber attacks

Large language models (LLMs) that generate text and chatbots that provide credible interactivity will emerge as powerful tools for cyber attacks. Social engineering attacks that have to date been conducted by human adversaries will become scalable to large populations and will still be personalized to individual users. For example, artificial intelligence will be able to conduct spearfishing attacks against entire communities, starting with customized emails or messages personalized based on browsing histories. Followed by interactive messages that adapt content based on the target’s responses. Detecting AI generated content will become a fundamental part of the cybersecurity arms race. Academic research is needed at a national scale to build the methods that make AI-generated misinformation and attacks robustly detectable, and also into the environments, interfaces, and tools to secure users' online interactions.

[8] See, e.g., Chapter 4 and Findings 4.12 and 4.13 of National Academies of Sciences, Engineering, and Medicine. 2022. *Cryptography and the Intelligence Community: The Future of Encryption*. Washington, DC: The National Academies Press. https://doi.org/10.17226/26168.

Rapid shift to hybrid work environments

The pandemic upended the workplace, significantly accelerating the trend to more remote work. The ability to keep society working even during a complete lockdown was to some extent a success of earlier investments in the digital ecosystem. We must now look to building a new class of tools to support hybrid workplace environments. Challenges include the urban, rural divide in terms of broadband access, equity issues across income and work types, education and issues related to privacy and security. Of course, as connectivity and remote work increases, so does the attack surface. Increasingly, as hybrid work environments become the norm, there needs to be a corresponding focus on cybersecurity in context of those hybrid work environments.

Opportunities and challenges for online education and training

This applies at all age levels. The challenges are acute for younger children. But perhaps new opportunities at high school and beyond levels, by bringing in new learning tools (AI agents), new materials, new set of teachers, etc.

Tools to to secure and improve our global supply chain

This issue has multiple components but advances in the digital ecosystem can help track goods, give early warning capabilities in terms of impending disruptions, and make it easier to re-engineer the network to move production and delivery of goods and services.

Tools for identifying, tracking and controlling the spread of (mis), (dis) information

Disinformation and misinformation targeted at individuals can have major impacts on system-wide responses to threats and challenges, as evidenced by vaccine hesitancy during the COVID-19 pandemic. With the deployment of AI based on large language and other generative models, the impact of disinformation on health, national security, and other aspects of society could become even greater. Identification and mitigation of disinformation is necessary, and individuals and communities need access to information that is contextualized and timely in order to make effective decisions.

New advances to collect, process and share personal information

Example areas that need these advancements in a manner that respects privacy, anonymity, fairness, etc. are mobility data, electronic health records, transaction data, and data collected by various body sensors. The COVID-19 pandemic showed the potential use of digital technologies such as contact tracing. Some countries used it effectively, others did not. The use of the technology is quite clear but requires further research into challenges and the potential misuse by individuals and organizations.

Biothreats and infectious disease surveillance

As there are progressive improvements in synthetic biology, biothreats will arise and need to be swiftly detected. More broadly detecting a multi-sector disruption that can potentially be engineered is needed. Opportunities for collecting, sharing information for infectious disease surveillance and monitoring (that includes environmental surveillance), and the ability to use new devices for sequencing are warranted. It is pivotal that these practices are connected in the large digital ecosystem that is secure by design.

Impacts of climate change

The increasing frequency of extreme events has the potential to stress and even de-stabilize the US power grids and, in turn, affect the communications networks that underpin the modern digital ecosystem. The security of this critical infrastructure is affected both by legacy systems that do not have protections built in and adoption of new technologies without proper protections. Federal support is needed to foster advances in computing research, together with interdisciplinary collaborations, to create robust, resilient next-generation infrastructure for the future. We need strategies for modeling the systems-level effects of climate change on the infrastructure that underpins our digital ecosystem, as well as decision-support systems and advanced algorithms to predict, identify, and mitigate failures (particularly cascading failures) in the context of those changes. In addition to leveraging decision-support systems to identify potential tipping points or security flaws in critical digital infrastructure, decision-support and visualization systems can be leveraged to optimize design of new infrastructure. Research in these areas will need to balance some important tradeoffs, as resiliency often requires redundancy and thus potentially reduces efficiency.

6. What further advancements to cybersecurity education and workforce development, at all levels of education, should be considered to prepare students, faculty, and the workforce in the next decade for emerging cybersecurity challenges, such as the implications of artificial intelligence, quantum computing, and the Internet of Things on cybersecurity?

In cybersecurity education and workforce development, there is often an assumption that people simply need to be taught how to use the technologies, rather than considering that the technology may simply not be useful for the user. Cybersecurity education and workforce development efforts - like research - should take a human-centered design approach that puts community stakeholders in the driver's seat. Promoting a Community Informatics approach to technology design, development, and buy-in may help address this challenge.

Experiential education (hackathons, capture the flag competitions, ethical hacking) has seen significant adoption and needs to continue.

Cybersecurity education for non-experts who form the vast majority of people who use and/or are impacted by the use of computers- policy, healthcare, etc. would help drive funding and research as more people would see the need for cybersecurity across sectors.

7. What other research and development strategies, plans, or activities, domestic or in other countries, should inform the U.S. Federal cybersecurity R&D strategic plan?

The R&D Goal related to quantum-resistant cryptography should remain but be updated in light of the progress in this area since 2019. In July, 2022, NIST announced[9] their first selections of quantum-resistant public-key encryption and digital signature algorithms for standardization, and NSA shortly followed with an update[10] to their Commercial National Security Algorithm suite (CNSA 2.0) incorporating some of the NIST selections. Further, in May 2022, President Biden issued National Security Memorandum NSM-10 ("Promoting United States Leadership in Quantum Computing While Mitigating Risks to Vulnerable Cryptographic Systems,"[11]) directing various activities within Executive Branch agencies to prepare them for the transition to quantum-resistant cryptography.

[9] https://nvlpubs.nist.gov/nistpubs/ir/2022/NIST.IR.8413-upd1.pdf
[10] https://media.defense.gov/2022/Sep/07/2003071834/-1/-1/0/CSA_CNSA_2.0_ALGORITHMS_.PDF
[11] https://irp.fas.org/offdocs/nsm/nsm-10.pdf

Also, in December 2022 the President signed P.L. 117-260, the “Quantum Computing Cybersecurity Preparedness Act”[12], which adds additional reporting requirements on civilian agencies. The R&D goals in the QIS section, specifically those related to quantum-resistant cryptography, should be revised to build on these recent government activities.

Both the National Defense Strategy[13] and the National Security Strategy[14] place significant emphasis on cybersecurity in topics ranging from norms of operating in cyberspace to resilience to cyber attacks. Coordination between non-defense and defense research and ensuring that advanced capabilities are transitioned into application and practice would be beneficial.

[12] https://www.congress.gov/bill/117th-congress/house-bill/7535
[13] https://www.defense.gov/National-Defense-Strategy/
[14]https://www.whitehouse.gov/wp-content/uploads/2022/10/Biden-Harris-Administrations-National-Security-Strategy-10.2022.pdf

APPENDIX

Here is the link to the CCC's Response to NITRD's "RFI on Update to the 2016 Federal Cybersecurity Research and Development Strategic Plan."

The following is a list of recent CCC RFI responses that have discussed issues which intersect with the issues raised by the RFI.

- Computing Community Consortium's Response to RFI "Request for Information on Advancing Privacy-Enhancing Technologies"
- Response to RFI on Federal Priorities for Information Integrity Research and Development
- The Computing Community Consortium's Response to Request for Information on Evaluating and Improving the NIST Cybersecurity Resources: The Cybersecurity Framework and Cybersecurity Supply Chain Risk Management
- Response to RFI on Public and Private Sector Uses of Biometric Technologies
- Computing Community Consortium (CCC) Response to Establishing Confidence in IoT Device Security: How do we get there?

The following is a list of CCC workshops since our 2019 response (associated community reports can be found at the link) that have discussed issues which intersect the issues raised by the RFI.

Artificial Intelligence

- Artificial Intelligence Roadmap Workshop 2 – Interaction
- Artificial Intelligence Roadmap Workshop 3 – Self Aware Learning
- Artificial Intelligence / Operations Research Workshop 1
- Artificial Intelligence / Operations Research Workshop 2

Post Quantum Cryptography

- Identifying Research Challenges in Post Quantum Cryptography Migration and Cryptographic Agility

Computing in Complex Human Systems

- Code 8.7: Using Computational Science and AI to End Modern Slavery
- CCC / Code 8.7 Workshop on Applying AI in the Fight Against Modern Slavery

Workforce/Economics

- Content Generation for Workforce Training
- Economics and Fairness

- The CCC Hybrid Workshop on Best Practices for Hybrid Workshops

Assured Autonomy

- Assured Autonomy Workshop #1
- Assured Autonomy Workshop #2
- Assured Autonomy #3

Health

- NAE/CCC Workshop on the Role of Robotics in Infectious Disease Crises
- Computational Support for Substance Use Disorder Prevention, Detection, Treatment, and Recovery

Climate

- Building Resilience to Climate Driven Extreme Events with Computing Innovations: A Convergence Accelerator Workshop

The following is a list of CCC whitepapers since our 2019 response (associated community reports can be found at the link) that have discussed issues which intersect the issues raised by the RFI.

Climate

- Computing Research for the Climate Crisis

Artificial Intelligence

- Imagine All the People: Citizen Science, Artificial Intelligence, and Computational Research
- Artificial Intelligence at the Edge
- Artificial Intelligence and Cooperation
- Interdisciplinary Approaches to Understanding Artificial Intelligence's Impact on Society
- The Rise of AI-Driven Simulators: Building a New Crystal Ball
- Next Wave Artificial Intelligence: Robust, Explainable, Adaptable, Ethical, and Accountable

Socio-Technical Computing

- An Agenda for Disinformation Research

- Modernizing Data Control: Making Personal Digital Data Mutually Beneficial for Citizens and Industry

Broad Computing

- A National Research Agenda for Intelligent Infrastructure: 2021 Update
- Pandemic Informatics: Preparation, Robustness and Resilience
- Infrastructure for Artificial Intelligence, Quantum and High Performance Computing
- Robotics Enabling the Workforce
- A Research Ecosystem for Secure Computing

Core Computer Science

- Post Quantum Cryptography: Readiness Challenges and the Approaching Storm
- Theoretical Computer Science: Foundations for an Algorithmic World
- Computing Research Challenges in Next Generation Wireless Networking