U.S. Leadership in Software and AI Engineering Workshop

Executive Summary

OCTOBER 2023

Carnegie Mellon University
Software Engineering Institute

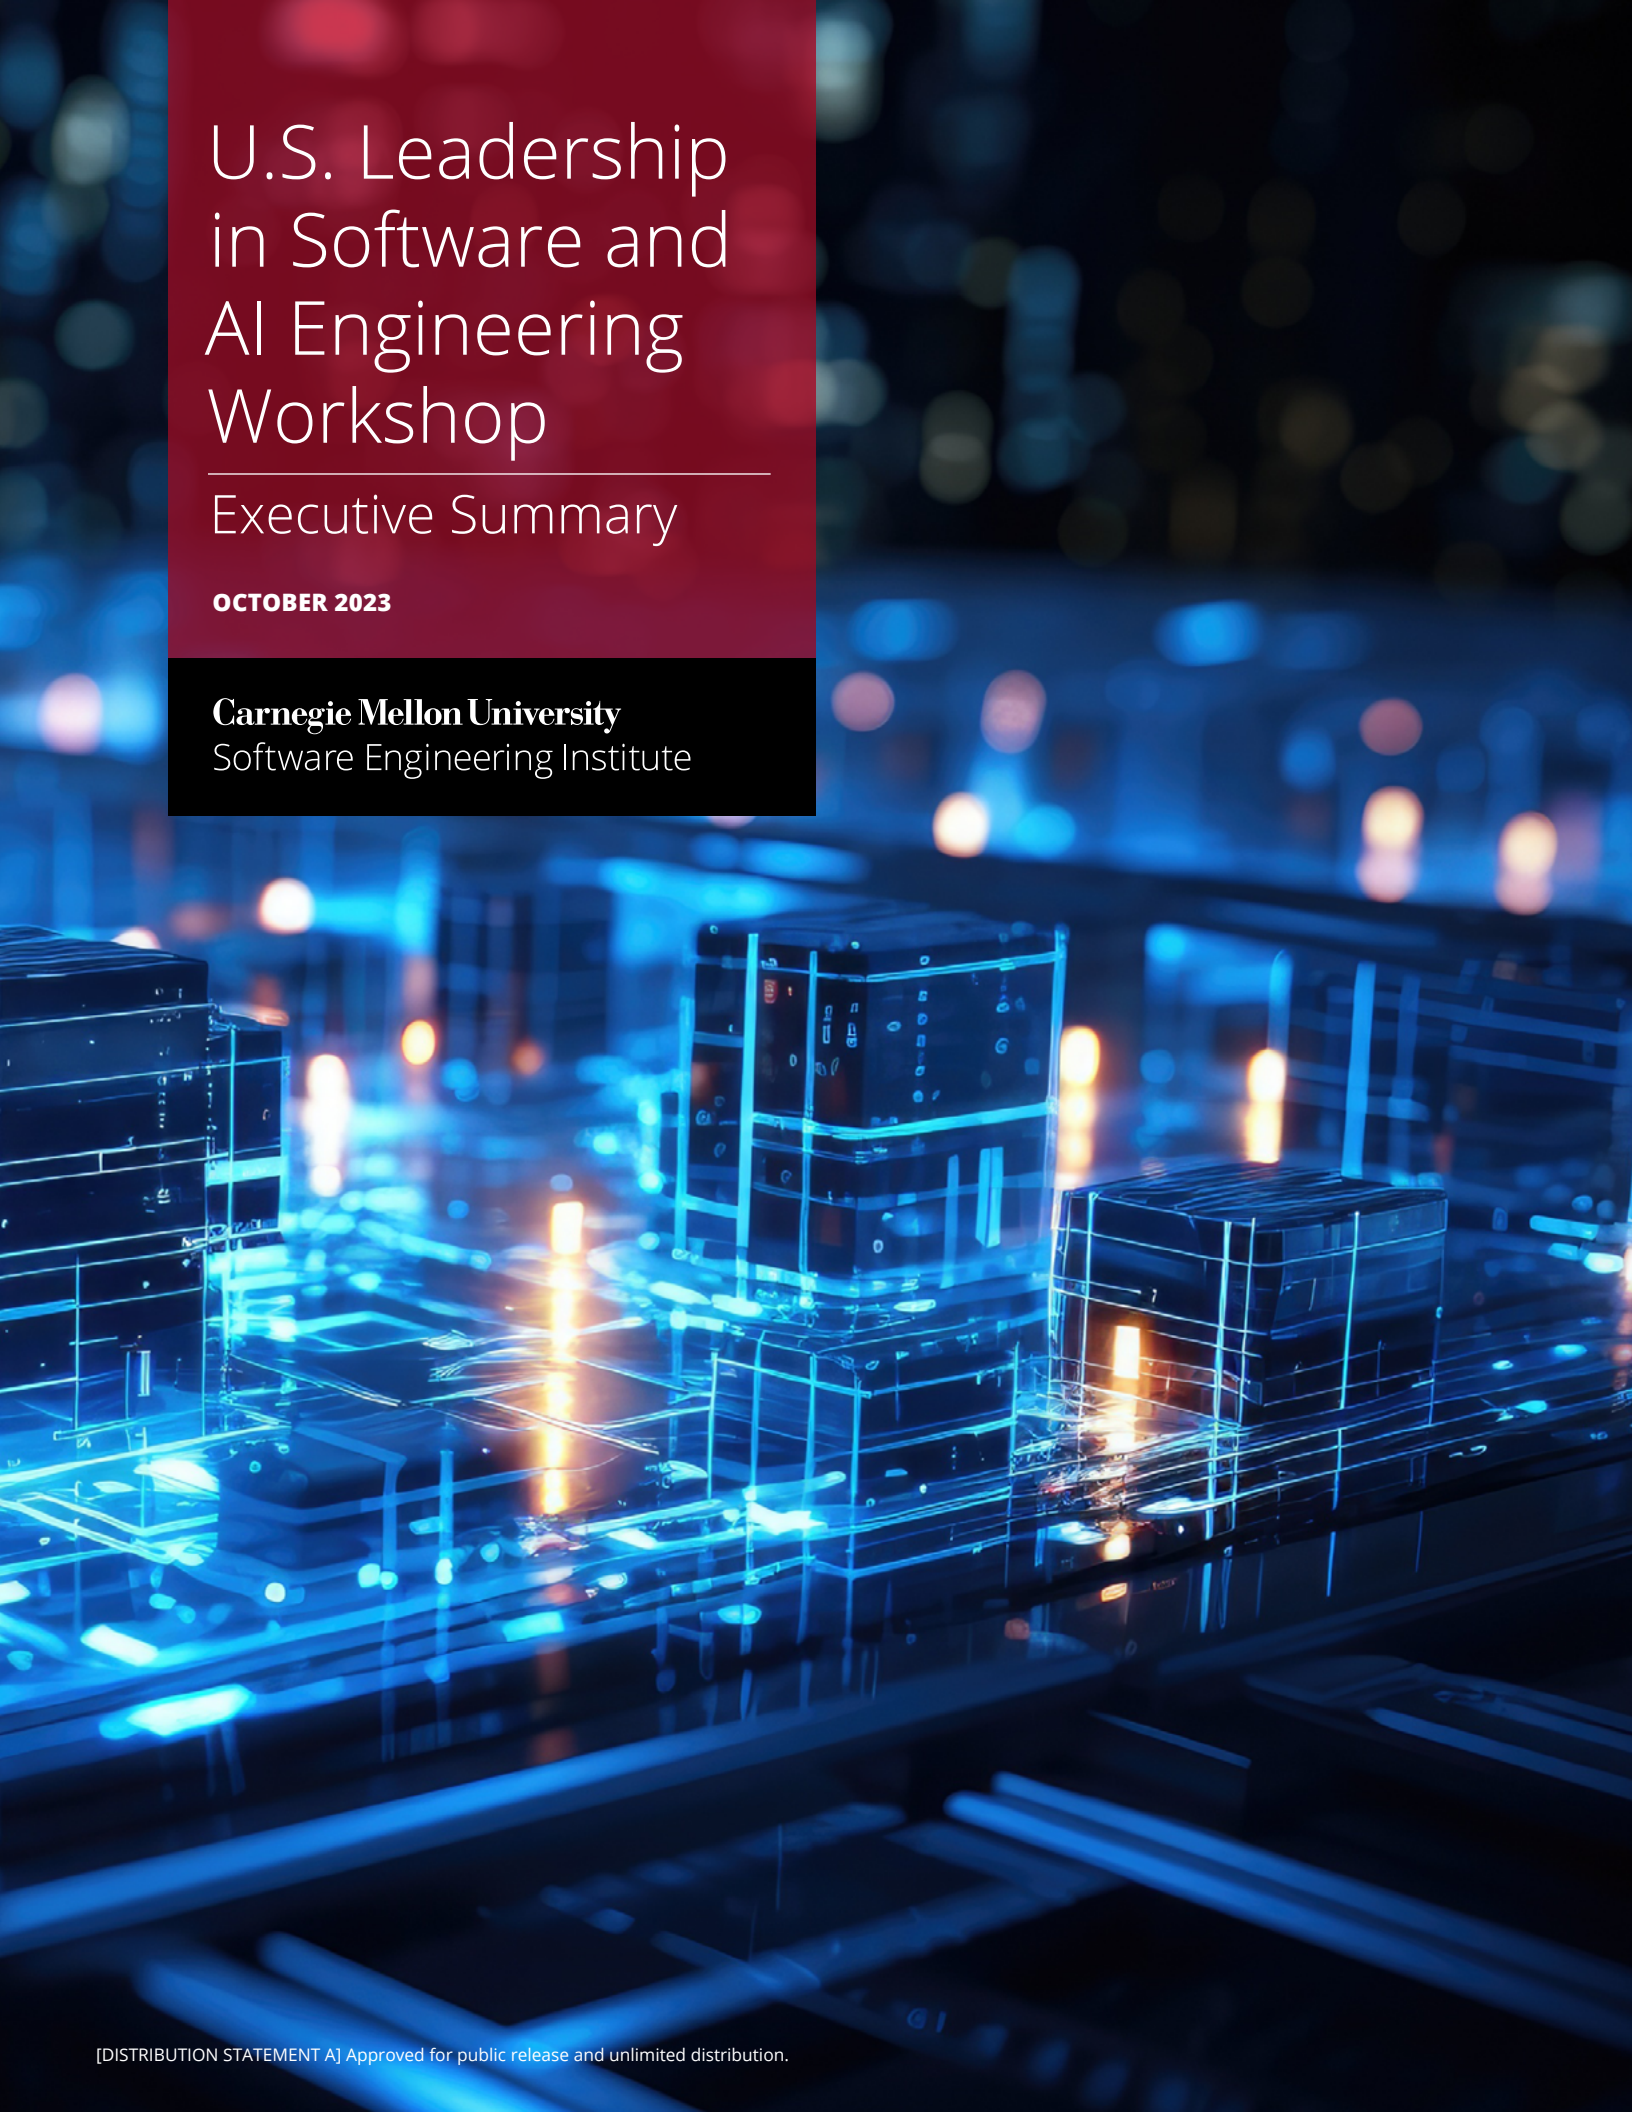

ADVANCES IN SOFTWARE ENGINEERING AND ARTIFICIAL INTELLIGENCE (AI) are providing critical and innovative capabilities across almost every domain, but the potential remains to do far more, particularly for applications that demand high levels of trustworthiness. To inform a community strategy for building and maintaining U.S. leadership in software engineering and AI engineering, the Software Engineering Institute (SEI) and the Networking and Information Technology Research and Development (NITRD) Program in the White House Office of Science and Technology Policy co-hosted a workshop at the National Science Foundation on June 20–21, 2023.

The event gathered thought leaders from federal research funding agencies, research laboratories, mission agencies, and commercial organizations to explore the fundamental research needed to support progress toward this goal. The workshop used the SEI's *Architecting the Future of Software Engineering: A National Agenda for Software Engineering Research and Development*[1] as a starting point because the areas of focus identified in the study have been confirmed as even more critical and urgent, particularly due to the rapid advances of generative AI in the two years since its release. Specifically, three research areas from the study were identified by participants as having direct relevance: AI-Augmented Software Development, Assuring Continuously Evolving Software Systems, and Engineering AI-Enabled Software Systems. Speakers and participants at the event worked to explore software-related challenges that are critical for multidisciplinary research across domains of importance to the nation as well as the promising research that is needed to engineer the necessary systems reliably and well.

WORKSHOP GOALS AND MOTIVATION

The workshop organizers brought together participants to encourage new partnerships that will advance U.S. leadership and national interests through the disciplines of software and AI engineering, and positively impact progress across virtually all scientific domains. Specific objectives for the workshop included

- Characterize how software engineering capabilities are having a direct impact on the future of our nation.
- Inform a community strategy for building and maintaining U.S. leadership in software engineering and AI engineering.
- Produce a report that summarizes challenges, opportunities, and strategic priorities.
- Identify research questions that energize the computing community and spark new collaborations.
- Identify updates to the Carnegie Mellon University (CMU) SEI *National Agenda for Software Engineering National Study and Roadmap.*

Executing and advancing the closely related disciplines of software engineering and AI engineering are indispensable to our ability to develop and deploy intelligent software systems effectively and rapidly. While the engineering of AI capabilities has unique and challenging requirements, these capabilities are implemented in software. To date, there has been significant research within software engineering on the technologies and practices needed to build such AI-enabled systems with confidence. While comparatively more recent, the fundamental theories, practices, and knowledge base for AI engineering are receiving significant research attention to ensure that AI capabilities are incorporated into systems with expected trustworthiness and responsibility.

There has also been considerable excitement around the idea of using AI to help in the engineering of software systems at scale. Approaches exploiting large language models (LLMs) are already automating some tasks that were thought to require human creativity, including some aspects of software engineering. As the boundaries of software and AI engineering blend, the tools and techniques available to engineers to develop top-priority capabilities are also changing. The rapidly changing technical environment creates further urgency to prioritize areas of most critical need and allocate multidisciplinary resources to the most challenging and essential areas of concern.

1 See **https://resources.sei.cmu.edu/library/asset-view.cfm?assetid=741193** to download a copy of the study.

CRITICAL NEEDS AND PRIORITIES: FIVE PRIMARY THEMES

In keynote speeches, breakout sessions, and lightning talks, participants almost unanimously remarked on the rapid acceleration of new technologies in the software development lifecycle and the role of AI in shaping the future of software systems. As the critical need for new approaches to navigate both the opportunities and the challenges was discussed, five main themes emerged.

1. AI is transforming the software engineering process and how we engineer software systems. The increasing symbiosis of humans and machines is transforming every phase of the software development lifecycle.

In software engineering, we are witnessing the emergence of a symbiotic workforce, where autonomous, intelligent assistants will work with software engineers to develop systems. This revolution in the way we approach software development will reshape the entire lifecycle, giving rise to approaches that promise to enhance productivity, quality, and efficiency. Software engineering should utilize AI tools and technology in the lifecycle, and software engineering principles should serve as a foundation for the development, evolution, and evaluation of AI-enabled software. The use of AI will likely make it possible to automate much harder programming and software-quality problems. While we recognize that tasks, skills, and tools will inevitably undergo transformation in this new paradigm, the specifics are not yet fully evident.

Current technological advances, especially those related to AI and machine learning (ML) tools, will fundamentally alter the ways in which applications are built—from design-to-code platforms and tools, to ML models that automatically generate code, to models that automate elements of application testing. ML-generated code is already in commercial codebases, and the overall percentage is already rapidly growing.

In fact, the experimental application of LLMs shows promise across the entire lifecycle. Effective application of LLMs may enable the ultimate "shift left" approach, where tasks that are traditionally done at a later stage of the process, such as testing or performance evaluation, can be done early, often before any code is written, or incorporated effectively throughout software development. Design-to-code platforms and tools could make it easier for developers to bring their ideas to fruition as models automatically generate code and streamline repetitive coding tasks. Leveraging advanced automation techniques, including AI- and LLM-enabled capabilities for everything from coding and code review to deployment at scale, integration test, and debugging, could streamline workflows, improve code quality, and accelerate the development cycle. Research exploring how to apply LLMs is only in its early phases, however, and many potential issues must be addressed, including the following:

- A substantial number of solutions have been trained on a single proprietary data source or on proprietary algorithms, and, as a result, it is not clear how robust their inferences and conclusions are.
- Filtering issues can make conclusions hard to replicate, especially since it is not always clear what kind of filtering has been done. Some models are trained on data that specifically omits some knowledge, and in other instances, the companies that own the models decide to censor some results.
- More diversity in models, systems, and applications is needed, and the research community should not put too much trust in a single model. Public funding might help address this issue by generating models and software/hardware infrastructures that remove the proprietary or black-box decision making that influences results.
- Given the speed with which innovations can be developed in this space, the software research community has become increasingly focused on quick prototypes as opposed to long-term, systematic research.
- Most effective techniques will likely be based on hybrid solutions, that is, a combination of LLMs, other AI, and data-driven automation approaches. Investigations of hybrid solutions should be accelerated.

While these new technologies promise to bring many benefits, they also have the potential to quickly multiply negative effects, such as security problems and AI debt (i.e., the cost of the complex mix of processes and procedures needed to discover, train, and deploy predictive models that are accurate and dependable). We need to develop sound and empirically based methods now for determining what approaches are considered successful and how to guide future software development lifecycle optimizations. Moreover, successful integration of AI in software development also relies on many non-technical factors, including the need for a "smart assistant" that understands team dynamics and roles and responds appropriately to human interactions and needs.

2. While generative AI has reached a level of sophistication that may seem to resemble human intelligence, it is considerably harder to determine the level of trust that should be placed in the outputs.

The assurance of mission- and safety-critical cyber-physical systems (CPS) has become increasingly challenging due to the growing complexity of these systems. The introduction of AI elements further compounds these difficulties because they can create large bodies of new code quickly, complicate the understanding of system behavior, and introduce new attack vectors, including the poisoning of training data and prompt injection, in which AI prompts can include code to generate pernicious behaviors.

As a result, while it is already clear that generative AI can make software developers more productive (in terms of producing code), there are well-founded worries about the quality and sustainability of the code produced. These new AI tools may already be producing a huge wave of technical debt that could overwhelm downstream software engineering efforts. In some studies, generative AI tools regurgitated old defects as often as they produced good fixes. Novice developers may lack the expertise to understand the limitations of the code being produced. AI-produced code will coexist alongside human-built code for a long time. We have few options to help end users and developers decide whether to trust code generated by tools and how this trust should compare to the trust in human-written code. Do we trust an AI tool more or less than a human, even if humans may make more mistakes? Where do we address trust: in the ML models themselves, in the software engineering, in testing, in how users interact with the system, or all of the above?

Research has already begun to identify the factors that can increase software developers' trust in AI tools. Key factors include source reputation, interaction (e.g., validation support and feedback loops), control (e.g., degree of ownership and autonomy), system features (e.g., ease of installation and performance measures), and expectations (e.g., how well the tool fits the style/goal of the developers). *Explainability* is not a proxy for *trustability*. By their nature, many AI systems cannot cogently explain why they arrive at their conclusions.

One goal should be increasing our ability to build trustable systems out of untrusted components. A second goal to explore is adopting AI to generate evidence about a resulting system that can be independently verified (e.g., analogous to the development of proof-carrying code or AI-generated code that comes with its own evidence). Another aspect of trust that requires research is whether AI tools leak intellectual property. It is possible that a model might learn on a proprietary codebase and then recommend pieces of that codebase to inappropriate users. Today we do not trust AI—but we do not always trust humans either. Rather than focusing on making AI trustworthy, we could use it to help us increase trust, using techniques such as generating evidence and incorporating AI into software testing and reviews.

Data assurance is another new frontier in the assurance of AI. In fact, it is one of the key components that makes assurance hard for AI, given the difficulty of understanding how data affects the final behavior of the system. The scalability of assurance for large AI models also poses a significant hurdle. Although some verification techniques have improved, the rapid increase in model size outpaces these approaches, which can render current verification methods inadequate from the outset.

3. **Redefining the discipline of software engineering to encompass the use of new technologies (including but not limited to generative AI), is imperative along with rethinking the associated curricula, tools, and technologies. This effort is key to designing and building, evolving, and evaluating trustworthy software systems in a responsible, ethical way.**

Redefining the software engineering discipline with AI is leading toward a revolution that changes how engineering solutions are explored, systems are built, and AI aids in the operation of systems. Education is a crucial aspect of any transformation effort brought about by AI, with new degrees and curricula incorporating AI into various engineering disciplines.

To keep up with the rapid advancement of AI technologies, software engineering curricula must include instruction on both the application of AI in the software engineering lifecycle and on how tools can facilitate the design, development, training, testing, and authorization of AI-enabled software. This evolution of software engineering curricula, both at the undergraduate and graduate levels, requires a dynamic component to ensure that the workforce is well equipped to effectively use these tools in supporting the development lifecycle.

Care must also be taken to make curricula equitable. Some initial observations as AI tools start to be used in software classes indicate that groups that are underrepresented in technology disciplines are also less comfortable using these technologies. This factor and others like it should be considered to avoid creating an environment where people with access to AI tools have clear advantages, and other groups without equitable access get left behind. Retaining talent in academia is also a concern. PhD students and faculty often face financial challenges due to the demanding nature of research and the need to secure funding. Efforts to make PhD programs more attractive, reduce funding restrictions, and provide sustained funding can help address these issues. The cost of an undergraduate education is also a significant concern. Government involvement in addressing the educational system's challenges can contribute to producing a workforce better equipped to address the nation's challenges effectively.

Enhancing fluidity between academia and other sectors can promote knowledge exchange. Incentivizing collaboration among universities and industry is crucial to address important research needs effectively. Key elements in fostering such collaboration include establishing public-modeled problems, data repositories, and testbeds to facilitate joint research efforts. Government agencies can also play a role by effectively utilizing commercial solutions and services where they prove beneficial and identifying bottlenecks that hinder progress.

4. New technologies, including generative AI, seem to hold the promise of making almost everyone a programmer. As a result, AI literacy and the development of new skills are needed throughout the workforce.

The landscape of programming is evolving dramatically. Instead of relying solely on those with traditional technical skills and expertise in software, systems, and AI engineering, new tools promise to enable almost everyone to become a "programmer." For this approach to be successful, new skills and abilities must be cultivated across a much wider range of people. These new skills and abilities include problem solving, critical thinking, and a general understanding of AI and ML.

The skills needed by professionally trained software programmers and engineers will also shift. While many traditional software engineering skills will likely become less valuable given AI tool capabilities, the value of the remaining skills may increase dramatically. For example, research results from Microsoft about its Copilot tool that generates code via LLMs indicate that users need to spend less time writing code but more time understanding and reasoning about code.

Software engineers will need a firm grasp of uncertainty and probabilistic reasoning, an increased capacity to detect problems and make informed design decisions, strong systems-thinking skills, and a keen awareness of the ethics of AI. The discipline of prompt engineering is beginning to gain traction, which involves programming in natural language and has potential applications in various stages of software development. Different prompts given to code models result in the generation of different code, highlighting the challenge of obtaining trustworthy output from these models.

Moreover, the potential impact on society and the economy of using AI in software systems necessitates that decision-makers and leaders in all domains comprehend the fundamental principles of AI and be competent in asking the critical questions to enable their trustworthy development and responsible use. Initiatives can be launched to provide training, workshops, and resources to ensure that individuals in positions of influence and authority are equipped to make informed decisions regarding AI technologies and their applications. By empowering leaders with AI literacy, we can foster the responsible and beneficial integration of AI in our lives.

5. The use of AI tools such as LLMs can mask the tradeoffs being made between the functionality of software systems and their safety and security. Research is needed to identify and make explicit the key engineering tradeoffs being made during the design, development, training, testing, and authorization of systems that include AI components.

Trust, trustworthiness, and confidence in software systems that include or are developed using AI components are top priority considerations. To achieve trustworthiness, engineers must navigate key tradeoffs in system development, ensuring the system performs as intended without overstepping its boundaries. This trust should extend as the system inevitably changes over time, providing measurable confidence in the system's evolving performance. Research is essential to

enable this outcome by providing mechanisms for identifying engineering tradeoffs throughout the specification, design, training, testing, and authorization of critical systems.

Explicit tradeoffs that set limits on AI systems are also needed to address concerns for both direct users and others potentially impacted by the system's actions or data. Although technologies like ChatGPT currently implement some features that prevent harm at the expense of performance, explicit engineering tradeoffs are needed during system development to clarify the relationship between functionality and safety/security. Research in AI-enabled systems must identify and analyze these tradeoffs explicitly to maintain safety and security throughout the software engineering lifecycle.

Additionally, AI-enabled tools should be designed to explicitly show the tradeoffs involved in developing a system instead of obfuscating or concealing them from key decision makers. Transparency in engineering tradeoffs is especially critical when incorporating technologies like smart coding assistants to ensure the development of robust and trustworthy systems.

RESEARCH NEEDS

Software and AI capabilities are advancing rapidly around the world and not just in high-resource nation states. They will continue to advance in complexity and sophistication without bound for the foreseeable future. To bolster U.S. leadership in this incredibly competitive domain, participants at the workshop identified a need to focus on research breakthroughs and development in software engineering and AI engineering, system architectures, and defining trustable systems. Presentations and discussions from multiple federal agencies showed the extent to which their plans for executing their missions rely on advanced software and AI capabilities.

Workshop participants also discussed the importance of improving collaboration mechanisms among academia, industry, and the federal space, including suggestions to invest in operationally relevant data sets and testbeds to enhance collaboration. Likewise, participants highlighted the need for open access to resources, such as models and data sets, in software engineering and the importance of breaking down large models into smaller pieces for better understanding and progress. The significance of social factors, access, and soft skills in AI and the importance of taking a multi-disciplinary approach were also acknowledged. The high-priority themes we identified also revealed a significant need for intentional crosscutting progress in data, standards, and all tradeoffs and aspects of trust. Specific areas of needed research discussed included:

- Software architectures for modern software needs. Architectures for AI-based systems should be developed so that they are resilient to attack and support federated data sources. The development of modeling and analysis techniques is needed to guide early design decisions, facilitate downstream test and evaluation (T&E), and enable evidence creation.
- AI engineering practices for trustworthy use of ML and LLM capabilities. Research is needed to enable the development of trustworthy systems to mitigate weaknesses in ML and LLMs and support ongoing updates to ML- and LLM-based capabilities as algorithms and training improve.
- Data-intensive software engineering. Software repositories have a wealth of information regarding current and older projects. There is a need to support repository mining for defect repair, API compliance, refactoring, synthesis, transformation, and evidence-based T&E. Data federation, privacy protection, and multi-institutional data collaboration are important challenges in integrating various types of data, such as health and environmental data.
- Diverse, advanced technical models and analyses to support development, evolution, and T&S. The use of modeling and analysis is essential in modern practice. Modeling and analysis must be integrated into practice in a way that allows a diversity of tools. More robust code models must be built by considering different code properties, such as syntax, semantics, and evolution, and incorporating them into the model's design and loss functions.
- Cybersecurity considerations for AI-reliant and software-reliant systems. Systems are growing in complexity and the number of interconnections, with larger external and internal attack surfaces, including AI attack surfaces. A focus on cyber risk is needed, including how to measure and manage attack surfaces, since threats are growing in sophistication and scale. Architectures devised for security and resiliency are needed as well as models and tools to enhance cybersecurity.
- Clear standards and guidance. There is a need for clarity in the development of standards for AI systems, as they are often asked to meet a large and varied number of requirements related to trustworthiness, security, privacy, and ethical considerations.

CONCLUSION AND NEXT STEPS

This workshop delved into various aspects of software and AI engineering, addressing challenges, opportunities, and ethical considerations. It highlighted the paradigm shift brought about by AI and LLMs, requiring alignment between models, researchers, and diverse user groups. Participants emphasized the need for transparency, trustworthiness, and collaboration across different sectors to effectively navigate the evolving landscape of AI technology.

The workshop also highlighted the impact of AI on various domains, including the workforce, cybersecurity, and autonomous systems, and the importance of collaboration and engagement with stakeholders was emphasized. The growing influence of AI in society, along with the acceleration of technology in general, demands interdisciplinary collaboration, technical advocacy for broader use cases, and policy development informed by the research community.

Making investment decisions in the right technical domains and fostering powerful partnerships is key to meeting the critical needs and priorities of the U.S. for software and AI engineering. For example, Figure 1 shows the actions taken to avoid the risks of a U.S. economy dependent on foreign chip manufacturing, which involves industry investments of around $50 billion and a proposed government investment of another $50 billion. AI technology investment followed a similar path, where a possible U.S. technology gap motivated major government and industry investment. The increasing awareness of the risks to national security and the U.S. economy motivated action in those cases, and those concerns also underscore the importance of making a similar strategic investment in software engineering research.

Chip Manufacturing

Risk: The U.S. economy has become dependent on foreign chip manufacturing.

- U.S. capacity fell to ~13% in 2015, compared to 30% in 1990 and 42% in 1980.
- 2020–2021: There were worldwide shortages related to COVID-19 pandemic.

U.S. Actions

- 2017: The President's Council of Advisors on Science and Technology (PCAST) reported on U.S. Leadership in Semiconductors.
- 2020–2021: Intel's $20 billion+ Taiwan Semiconductor Manufacturing Company represented $30 billion+ in U.S. fabrication investments.
- 2022: The CHIPS Act was signed into law, including $52.7 billion for American semiconductor research, development, manufacturing, and workforce development.

AI Technology

Risk: The U.S. AI technology gap compares negatively to other nation states.

- Many nations are interested, but it is primarily a two-nation race.
- Multiple nations are announcing multi-billion-dollar investments in AI.

U.S. Actions

- 2018: The DARPA "AI Next" Campaign invested $2 billion.
- 2019: The *Executive Order on AI* was released.
- 2021: NITRD investment #1 of 12 was made.
- 2021: The National Artificial Intelligence Initiative (NAII) was established through bipartisan legislation.
- 2023: The White House announced a $140 million investment to create seven AI research hubs.

Software Engineering Research

Risk: Software engineering advances have not kept up with the critical nature of software for U.S. national security and competitiveness.

This is important because

- Software is the backbone of critical systems.
- Software includes complex supply chains.
- Software is infrastructure.

Initial U.S. Actions

- 2019–2020: The NITRD Future Computing Community of Interest; *National Strategic Computing Initiative Update;* and Software Productivity, Sustainability, and Quality Working Group were formed.
- 2021: CMU SEI's *A National Agenda for Software Engineering Research and Development* study was published.
- 2023: U.S. Leadership in Software Engineering and AI Engineering: Critical Needs & Priorities Workshop was held.

Figure 1. Landscape of U.S. Investment in Critical Technologies

About the SEI

Always focused on the future, the Software Engineering Institute (SEI) advances software as a strategic advantage for national security. We lead research and direct transition of software engineering, cybersecurity, and artificial intelligence technologies at the intersection of academia, industry, and government. We serve the nation as a federally funded research and development center (FFRDC) sponsored by the U.S. Department of Defense (DoD) and are based at Carnegie Mellon University, a global research university annually rated among the best for its programs in computer science and engineering.

Contact Us

CARNEGIE MELLON UNIVERSITY
SOFTWARE ENGINEERING INSTITUTE
4500 FIFTH AVENUE; PITTSBURGH, PA 15213-2612

sei.cmu.edu
412.268.5800 | 888.201.4479
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.

DM23-0890