34th Annual INCOSE
international symposium
hybrid event
Dublin, Ireland
July 2 - 6, 2024

The Updated SERC AI and Autonomy Roadmap 2023

Tom McDermott
Stevens Institute of Technology
Hoboken, New Jersey, USA
+1-678-522-6569
tmcdermo@stevens.edu

Kara Pepe
Stevens Institute of Technology
Hoboken, New Jersey, USA
kpepe@stevens.edu

Megan Clifford
Stevens Institute of Technology
Hoboken, New Jersey, USA
mcliffor@stevens.edu

Abstract. The first Systems Engineering Research Center (SERC) Artificial Intelligence (AI) and Autonomy Research Roadmap was developed in 2020 and published in the first quarter 2021 special INSIGHT issue on Systems Engineering and AI. This roadmap development was heavily informed by the INCOSE Future of Systems Engineering (FuSE) initiatives. Following on in 2020, INCOSE and the SERC collaborated with the Association for Advancement of AI (AAAI) to execute two workshops entitled "AI meets Systems Engineering." These resulted in version two of the roadmap which was published as an introductory chapter to the book "Systems Engineering and Artificial Intelligence." In 2020 through 2023 the SERC hosted four SE4AI/AI4SE workshops with the U.S. Army that have further informed research and application at the intersection of AI and SE. This paper presents the updated version of the roadmap resulting from engagement across those four workshops. It is provided as a means to inform the SE community of the critical research needs and related applications emerging at the intersection of AI and SE.

Keywords. Artificial Intelligence, Machine Learning, Autonomy, AI4SE, SE4AI, FuSE

Introduction

In 2019, the Research Council of the Systems Engineering Research Center (SERC), a US Defense Department sponsored University Affiliated Research Center (UARC), developed a set of roadmaps (SERC, 2019) structuring and guiding four areas of systems engineering research: digital engineering, velocity, security, and artificial intelligence (AI) and autonomy. In its strategic research planning, the SERC employs a Research Council composed of senior faculty across our university collaboration network. The use of a graphical roadmap (as opposed to a text document) was selected as a way to make our research strategy more accessible to our sponsors and the larger systems engineering community. The SERC's current research strategy now aligns to the four core research areas supported by the four cross-cutting research strategies, as shown in Figure 1.

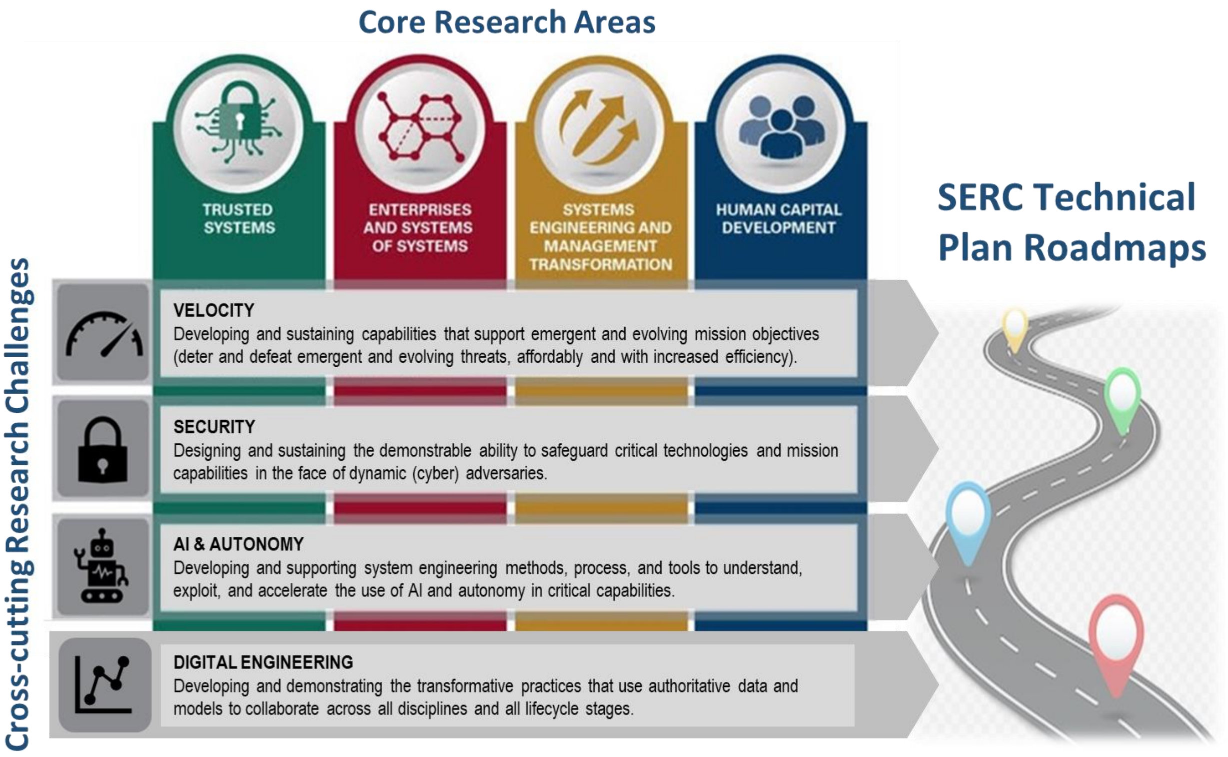
Core Research Areas
TRUSTED SYSTEMS
ENTERPRISES AND SYSTEMS OF SYSTEMS
SYSTEMS ENGINEERING AND MANAGEMENT TRANSFORMATION
HUMAN CAPITAL DEVELOPMENT
SERC Technical Plan Roadmaps
Cross-cutting Research Challenges
VELOCITY
Developing and sustaining capabilities that support emergent and evolving mission objectives (deter and defeat emergent and evolving threats, affordably and with increased efficiency).
SECURITY
Designing and sustaining the demonstrable ability to safeguard critical technologies and mission capabilities in the face of dynamic (cyber) adversaries.
AI & AUTONOMY
Developing and supporting system engineering methods, process, and tools to understand, exploit, and accelerate the use of AI and autonomy in critical capabilities.
DIGITAL ENGINEERING
Developing and demonstrating the transformative practices that use authoritative data and models to collaborate across all disciplines and all lifecycle stages.

Figure 1. SERC Research Areas and Roadmaps.

History of the AI and Autonomy Roadmap

Initial research needs for AI & Autonomy were discussed at a SERC Research Council meeting in December 2018. These included concepts of AI and machine learning (ML) capabilities in SE as a means to manage complexity and build trust, evolution of SE practice to adapt to emergent autonomous system behaviors, the opportunity with digital engineering to blend system knowledge representation and AI/ML, and related evolution of SE knowledge and skills. A Future of Systems Engineering (FuSE) workshop at the INCOSE International Workshop in January 2019 (IW2019) defined some of the initial vision and interaction for SE and AI, and first noted the application areas “SE4AI” and “AI4SE” as transformational opportunities for SE practice. The INCOSE AI working group was initially formed in 2019 and created a set of relationships between INCOSE, the SERC, and the Association for Advancement of AI (AAAI) to survey ongoing research and explore research needs. Version 1 of the SERC AI and Autonomy research roadmap was published in the first quarter 2020 special INSIGHT issue on Systems Engineering and AI (INCOSE, 2020). That version identified four research areas: AI/ML Technology, Automation and Manned-Unmanned Teaming, Digital Engineering, and SE Process Evolution.

In 2020, INCOSE and the SERC collaborated with the Association for Advancement of AI (AAAI) to execute two workshops entitled “AI meets Systems Engineering.” These workshops explored both the explosion of machine learning (ML) applications happening today, and how these will evolve to more dynamic human-machine interactions in teams. The AAAI interaction resulted in version 2 of the roadmap which was published as an introductory chapter to the book “Systems Engineering and Artificial Intelligence” (Lawless et al. 2021). That version changed “SE Process Evolution” to “Augmented Engineering,” recognizing the “AI4SE” would primarily be used to augment decisions of SE practitioners. This version also added a fifth research area on Workforce and Culture. These roadmap versions are not discussed further in this update, but are organized in a similar form to this discussion in those publications.

In 2020 through 2023 the SERC hosted a series of four SE4AI/AI4SE workshops with the U.S. Army that have further informed research and application at the intersection of AI and SE (SERC 2020, SERC 2021, SERC 2022, SERC 2023, SERC and INCOSE 2023)). This paper presents the updated version of the roadmap resulting from engagement and research results presented across those four workshops. It is provided as a means to inform the SE community of the critical research needs and related applications emerging at the intersection of AI and SE. It is our hope by sharing this work we will guide not only SERC research but also the transformation of the systems engineering discipline in general.

The AI and Autonomy Roadmap 2023

Each roadmap has a set of research "vectors" (arrows in the roadmap diagrams) leading to a visionary outcome or set of outcomes, and a set of capabilities (dots in the roadmap diagrams) we believe are needed to meet those long term outcomes. The listed capabilities in these roadmaps reflect not only SERC research, but other areas of research either known to be active or prioritized by our sponsors and the systems engineering community in general and our sponsors.

The current AI/Autonomy roadmap is shown in Figure 2. The envisioned long-term outcome of the SERC AI and Automation roadmap is "Human-Machine Co-learning." This outcome captures a future where both humans and machines will adapt their behavior over time by learning from each other or alongside each other. More importantly for systems engineering, this is a lifecycle model that is not envisioned and supported by most of the current-day systems engineering practices.

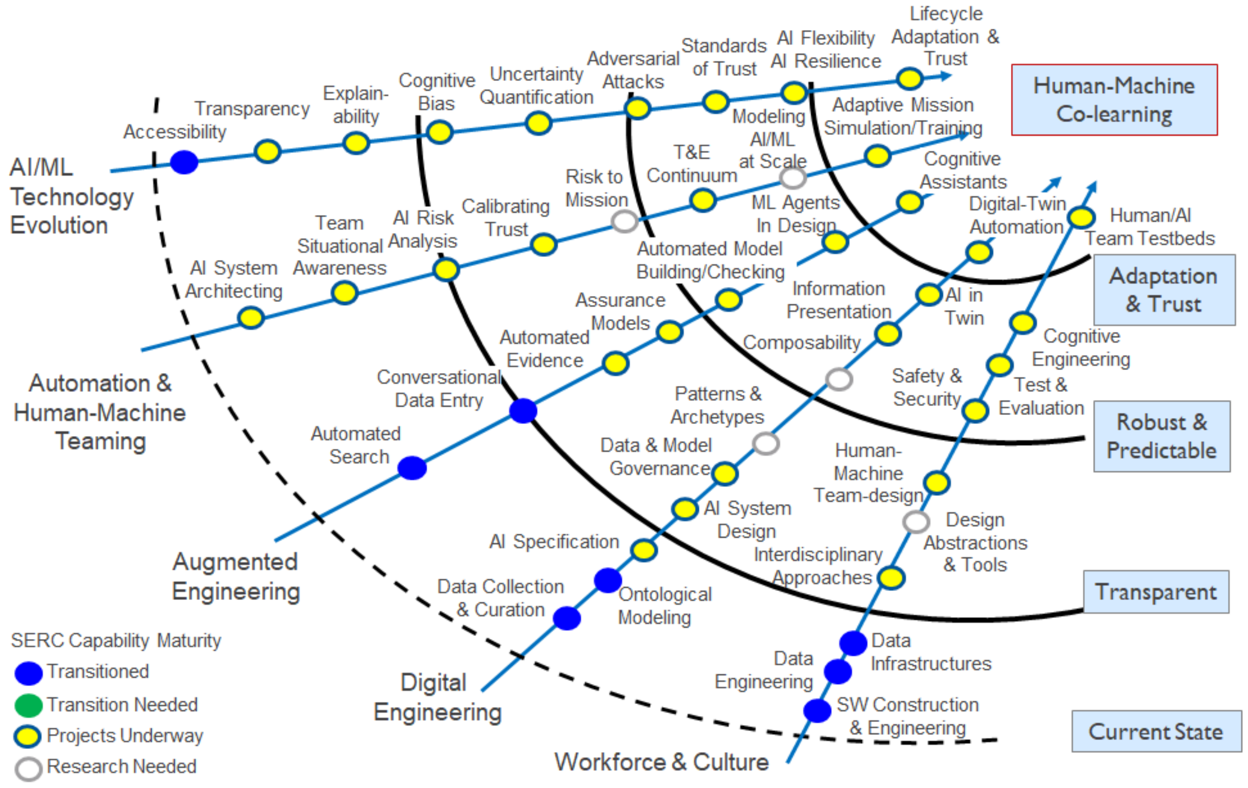

Figure 2. Current SERC AI and Autonomy Roadmap.

To achieve this end state, one might consider there is a need for both the AI and SE disciplines to pass through a set of "waves" or eras. Each of these defines a set of quality attributes for systems that employ human-machine co-learning. The first of these includes sets of technologies and approaches that make the decisions produced by AI systems more transparent to the human developers and users. The second wave is to produce systems that learn but are also appropriately robust and predictable in the type of critical applications normal to SE. The third wave involves adaptation and trust: systems that actually adapt and learn dynamically from their environments, and how humans develop trust in these systems.

The vectors of this research roadmap span five categories. The first of these vectors, AI/ML Technology Evolution, recognizes that the technological implementation of AI systems will evolve and will need to evolve in directions relevant to SE. Most of these can be related to the development of transparency and trust in technology. The second vector, Automation and Human & Machine Teaming, recognizes that the use of AI in systems will generally provide automation of human tasks and decisions and must be trusted. The third vector, Augmented Engineering, recognizes that AI technologies will gradually be used more and more to augment the work of engineering, what we call AI4SE. The fourth vector, Digital Engineering, recognizes that the current digital engineering transformation will be enabler for the use of AI in SE. The final vector recognizes a transformation will need to be accomplished in SE Workforce & Culture, with significantly more integration of software engineering and human systems engineering toward the forefront of SE.

The following sections walk through each of the research areas defined for each vector (these are the colored circles in the roadmap. Circles that are blue indicate research that has already transitioned to practice at the time this roadmap version was published. Green circles indicate research that is ready to transition. Yellow circles indicate research that is currently underway. White circles are research areas where (in our estimation) research is needed but not yet started. Again, these circles represent research in the SERC network of collaborating universities as well as in other organizations that we are aware of. An interesting aspect of this roadmap is the amount of research currently underway (yellow circles) and the lack of transition into the SE domains (blue circles). The next sections provide a short description of each research area.

AI/ML Technology Evolution

The algorithms that underly AI and ML are generally mature technology and are widely accessible through both commercial and open-source software tools. Accessibility of these algorithms, along with increases in computation and data storage, is driving the explosion of ML technologies today. We foresee that AI algorithms and methods will continue to become more available in tools that can be used by multiple disciplines. As systems engineers, we are always interested in the rigor with which these applications will be developed and tested. There are a set of continuing technology related challenges that SE practitioners would call “quality attributes” which remain significant areas of research. The SE discipline should be highly invested in exploring these qualities, others may not. Research areas include:

- **Accessibility**: AI/ML algorithms and methods will become more available in tools that can be used by multiple disciplines. For example, tools like ChatGPT that create accessibility to large language models (LLM’s) are beginning to emerge as chat interfaces to traditional SE tools.
- **Transparency**: For rigor and trust, SE techniques are needed for training AI algorithms that are created without hiding the computation behind the algorithms decisions. This research area was added to the 2023 roadmap, recognizing that each AI/ML algorithm has its own system lifecycle, and we lack rigorous SE processes for managing lifecycles at that level.
- **Explainability**: This research area relates to the development of sets of machine learning technologies and techniques that produce more explainable models, while maintaining a high level of learning performance (prediction accuracy); and enable human users to understand, appropriately trust, and effectively manage the resulting automation. Transparency and Explainability represent process and technology means to create more trustworthy AI/ML.

- **Cognitive Bias**: Systems approaches are needed for countering intentionally or unintentionally misleading decision-making in AI systems, caused by the training data, adaptative learning over time, or intentionally misleading information.
- **Uncertainty Quantification**: Systems approaches are needed for representing the uncertainty of AI prediction algorithms as well as the sources of uncertainty.
- **Standards of Trust**: Research on the nature of trust in AI/ML systems is needed to extend engineering standards of trust (dependability, availability, etc.) to also include socio-technical aspects of learning-based systems (ethics, understanding, validation, etc.). This research area was also added to the 2023 roadmap, recognizing that new system qualities are emerging in learning-based systems that have not been traditionally addressed in the SE and other engineering disciplines.
- **Adversarial Attacks**: AI/ML system security engineering must address new classes of attacks and security design methods. These include the use of adversarial samples to fool machine learning algorithms; defensive techniques for detection/classification of adversarial samples; and security of AI systems.
- **AI Flexibility and Resilience**: Research is needed on new SE methods to address operational resilience of the system and its flexibility to users incorporating AI, particularly involving the characteristics of ML systems at the boundary to human operators.
- **Lifecycle Adaptation and Trust**: The end-state of this research vector recognizes AI performance and related aspects of human trust will evolve over the lifecycle of a system as the system changes/evolves. The concept of trust in AI has been added to this 2023 version of the roadmap. SE must broadly address methods and practices to manage trustworthiness of AI/ML systems and technologies across all its practices.

Automation & Human-Machine Teaming

There is a lack of testbed environments in the research community to explore the end-effects of human machine interactive teaming. Many of the collaborative behaviors are not well understood when first developed and learning is required on both the engineering and user side to evolve effectiveness. Building appropriate data and live and virtual system architectures to support learning and adaptation is a critical research area. More agile change processes are also critical. Methods, processes, and tools are needed to connect system risk analysis results with AI software modules related to those risks. This is very similar to the cyber resilience research area. AI systems that self-adapt while maintaining rigorous safety, security, and policy constraints are not widespread today, so this is a significant research area. Methods for addressing AI-related system test and evaluation, particularly when these systems' ability to adapt and learn from changing deployment contexts improves. One largely unexplored area is AI/ML at Scale: appreciation for the dependence of an AI's outputs on its inputs. Scale in AI-based systems will increasingly lead to more general intelligence and an inability to relegate AI to a particular subsystem or component – in other words the problem becomes difficult to decompose. Computer-based simulation and training supporting non-static objectives and/or goals (games, course of action analysis) necessary to provide contextual learning environments for these systems.

- **AI System Architecting**: New SE practices are needed for building appropriate data and live and virtual system architectures to support learning and adaptation. An emerging research area added to the 2023 roadmap is the liv and virtual integration of joint human and machine learning, and related modeling and simulation practices.
- **Team Situational Awareness**: Methodologies are needed for supporting individual and team situation awareness extended to human-machine teams in multi-domain operations, primarily information modeling and interface design. This research area was added to the 2023 roadmap, and is further described in the next section.

- **AI Risk Analysis**: Methods, processes, and tools are needed to connect system risk analysis results with AI software modules related to those risks. An example might be an “AI Readiness Level” metric.
- **Calibrating Trust**: AI systems in critical applications must self-adapt while maintaining rigorous safety, security, and policy constraints. Research is needed on the SE methods and tools to constrain these systems, and the analytical and evaluation methods to assess these constraints.
- **Risk to Mission**: Trust in AI/ML systems arises from human and AI collaboration in shared tasks and functions, often arising from errors and unforeseen conditions. Research is needed to relate AI/ML technology performance and risk factors to mission level outcomes, and to create simulation platforms that effectively address these risks. The area was also added to the 2023 roadmap.
- **T&E Continuum**: Methods are needed for addressing AI-related system test and evaluation as a continuum of activities, addressing these systems’ ability to adapt and learn from changing deployment contexts.
- **Modeling AI/ML at Scale**: Deploying multiple AI/ML algorithms in multiple learning systems at systems of systems (SoS) scales requires evolving and new SE and SoS practices. Appreciation for the dependence of an AI’s outputs on its inputs; scale in AI-based systems will increasingly lead to more general intelligence and an inability to decompose and relegate AI to a particular subsystem or component. Some of these trends are starting to become prevalent in the AI algorithms used in self-driving vehicle modes.
- **Adaptive Mission Simulation and Training**: Computer-based simulation and training supporting non-static objectives and/or goals (games, course of action analysis) is necessary to provide contextual learning environments for these systems, leading to a “train as you fight” paradigm for AI’s.

Augmented Engineering

AI and ML have significant potential to help engineers and especially systems engineers do their work. We call this Augmented engineering. Automated search algorithms will be very beneficial, applying ML to historical data and relationships in the engineering domains. Human/computer interaction processes to convert natural language and other media to formal models should follow. Research is growing on automated construction of models from features in semantic data, used in both creation of new models and correctness of developed models. Automation of certification and accreditation processes via models and automation of quality assurance data will improve the reliability of future systems. This includes automation of evidence-based models for assuring correctness and completeness of system requirements and design. Research is underway on Cognitive Digital Assistants - conversational systems automating many mundane data entry, exploration, and engineering calculation tasks, and many workflows.

- **Automated Search**: Applying ML to historical data and relationships in the engineering domains can improve efficiency and reduce cycle times for engineering decisions and design. SE tools are already deploying automated search features using LLM’s and other AI algorithms.
- **Conversational Data Entry**: Human/computer interaction processes to convert natural language and other media to formal models will rapidly change the nature of SE tools. This is another area where the use of LLM’s is rapidly improving SE methods and tools.
- **Automated Evidence**: This research area and the Assurance Models research area address the need for automation of system assurance cases and related data given the complexity of modern systems. Research is needed to produce methods and tools for automation of certification and accreditation processes via digital models and related automation of quality assurance data. DARPA’s Automated Rapid Certification of Software (ARCOS) program is an example research effort.
- **Assurance Models**: Research is needed to develop methods and tools for automation of evidence-based models for assuring correctness and completeness of system requirements and design.

- **Automated Model Building/Checking**: Automated construction of models from features in semantic data is another emerging capability of LLMs, used in both creation of new models and evaluating correctness of developed models. This will rapidly improve as foundational models evolve to include more graphical capabilities and mathematical relationships.
- **ML Agents in Design**: Most simulations have inherently predefined outcomes. Embedding ML agents into simulations is needed to predict and evaluate new emergent behaviors in learning-based systems. This research area was added to the 2023 roadmap.
- **Cognitive Assistants**: This research vector culminates in the concept of digital cognitive assistants for various SE and engineering and related program management tasks – "augmented engineering". Cognitive assistants are Conversational systems automating many mundane data entry, exploration, and engineering calculation tasks, and many workflows.

Digital Engineering

Digital engineering will be a great enabler for use of AI/ML into engineering functions. Many of these research areas were mentioned in the Digital Engineering roadmap. Systems engineering will need to manage specific activities to build infrastructure and collect and manage data needed for engineering and programmatic activities in system development and support. As mentioned earlier, ontological modeling of engineering and programmatic data providing interoperability through standard and domain specific ontologies will be critical. Lifecycle management, control, preservation and enhancement of models and associated data will be a core SE activity to ensure value for current and future use, as well as repurposing beyond initial purpose and context. In the long-term AI should enable digital twin automation: fully dynamic virtual system copies built from the same models as the real systems running in parallel to physical systems and updating from the same data feeds as their real counterparts.

- **Data Collection & Curation**: Specific activities to build infrastructure and collect and manage data needed for engineering and programmatic activities in system development and support must be defined in SE processes.
- **Ontological Modeling**: Research is needed on data representation and modeling in the engineering and program management domains to support development of AI/ML augmented engineering tools. General research is needed for knowledge representation of engineering and programmatic data providing interoperability through standard and domain specific ontologies.
- **AI Specification**: Research is needed to define system-level and formal specifications for AI behaviors supporting subsequent verification activities.
- **System Design for AI Performance**: System design approaches are needed as a mechanism for generalization of AI/ML performance factored into design activities. This research area was added to the 2023 roadmap.
- **Data & Model Governance**: This research area covers lifecycle management, control, preservation and enhancement of models and associated data to ensure value for current and future use, as well as repurposing beyond initial purpose and context.
- **Patterns & Archetypes**: Widely used modeling constructs for AI/ML behavior and performance must be developed that separate design from implementation, supporting better reuse and composition. These should be well recognized patterns and should be supported by model libraries.
- **Composability**: As patterns and libraries become more widely utilized, rapid development and integration of design using higher level abstracted components and patterns will emerge, across multiple disciplines.
- **Information Presentation**: Research is needed in the area of AI-Vis: visualization approaches and interfaces supporting human-machine real-time collaborative information sharing via multiple media.

- **AI in the Digital Twin**: New uses of AI in digital twins enabling new functional and performance value will emerge from the development of system level digital twins. This research area was added to the 2023 roadmap.
- **Digital Twin Automation**: This is the expected outcome of digital engineering and AI: fully dynamic virtual system copies built from the same models as the real systems running in parallel to physical systems and updating from the same data feeds as their real counterparts. Research is needed to both support engineering of digital twins and to evolve future use cases.

Workforce & Culture

AI and Autonomy in the engineering domain, particularly the physics-based disciplines, requires much more interdisciplinary learning of data infrastructures, data engineering and software construction and engineering than is typically taught today. AI systems are highly interdisciplinary. Human-machine teaming is also very interdisciplinary, requiring knowledge across disciplines of machine control, cognitive science, and human learning. The traditional SE specialty disciplines such as safety and security must adapt their practices to non-deterministic processes and systems. Test and evaluation must be integral to development and continually evaluating the system. All of these dimensions will create workforce challenges for system developers and multidisciplinary challenges for educators.

- **Software Construction and Engineering**: AI is software and systems engineers need to understand the methods, practices, and tools for constructing AI and data-analytics software. The SERC developed the Digital Engineering Competency Model (DECM) to outline these base competencies.
- **Data Engineering**: Engineering of systems of data will become integral to and equal to engineering of the systems themselves. These practices are evolved in the software and data analytics realm but need to be more prevalent in SE and engineering roles.
- **Data Infrastructures**: The underlying data and data management activities must become a core systems engineering lifecycle process area. This process area has not been documented yet in SE lifecycle models.
- **Interdisciplinary Approaches**: Systems management of AI as a technical discipline with the other engineering and management related disciplines is an issue because the educational domains are separate. SE must become a translator across these domains with associated processes and decision tools.
- **Design Abstractions and Tools**: Increasing the level of abstraction of ML and AI methods and tools will improve application by other disciplines (again LLM’s are a current example). Just as “prompt engineering” has become a new skill for LLM’s, other similar skills will continue to evolve as AI tools advance.
- **Human-Machine Team Design**: New multidisciplinary approaches for the design of human-AI teams that follow best practices in Human Systems Integration (HSI) are needed, particularly those that relate human and machine tasks to functional design and then to requirements analysis. HSI will move more to the forefront of SE practice.
- **Safety and Security**: Skillsets, tools, methodologies, and policies related to safety, security, reliability and resilience of systems with AI must be developed and taught.
- **Test and Evaluation**: New methods and tools focused on adaptive T&E of learning-based and probabilistic AI systems must be developed and taught.
- **Cognitive Engineering**: Research is needed for more systematic use of cognitive, social, and experimental psychology to design and develop engineering systems to support the cognitive processes of both humans and machines.

- **Human-AI Team Testbeds**: Flexible testbeds for evaluating human-AI teams are needed for research, evaluation, design optimization, and adaptation. This research area was added to the 2023 roadmap.

Human-AI Teaming

In 2022 the U.S. National Academies of Science, Engineering, and Medicine released their consensus study on Human-AI Teaming: State-of-the-Art and Research Needs (NAP 2022). This provided the opportunity to review the SERC and National Academies of Sciences (NAS) research roadmaps and update the SERC AI and Autonomy roadmap based on critical research needs in that report. This report stated a primary concern that humans have, continue to be, and will further be "challenged in performing as successful monitors of complex automation, including AI systems" (NAP 2022). Most of the research areas in the SERC AI and Autonomy Roadmap also involve concerns related to relationship between the human and the machine automation, both beneficial and potentially harmful. This report is essential reading for SE's as this is the core of our SE challenge amidst the explosion of AI/ML applications in the current state. The NAS report identified research areas across nine objectives (NAP 2022):

1. Models and metrics for human-AI teams, considering the human and AI as a team
2. Improving team processes as a designed system, reflecting both team tasks and teamwork
3. Methods for improving situational awareness of AI systems, including AI awareness of humans
4. AI transparency and explainability
5. Human-AI interaction approaches, including roles and flexibility in levels of automation
6. Measures of trust that draw on human cooperation as opposed to machine dependability
7. Reduction of both human and AI bias
8. Approaches for training human-AI teams and related systems and platforms
9. Advances in the foundations of human-systems integration (HSI) processes, as well as SE processes

As a response to this report, we reevaluated the SERC roadmap to define the synergies between SERC identified research areas and those in the NAS report. We found there is a significant research need associated with the concept of AI System Architecting: how long-term, distributed, and agile human-AI teams will adapt through improved team assembly, goal alignment, communication, coordination, social intelligence, and the development of a new human-AI language. Figure 3 shows these relationships between the existing SERC research areas and the concepts related to AI Systems Architecting.

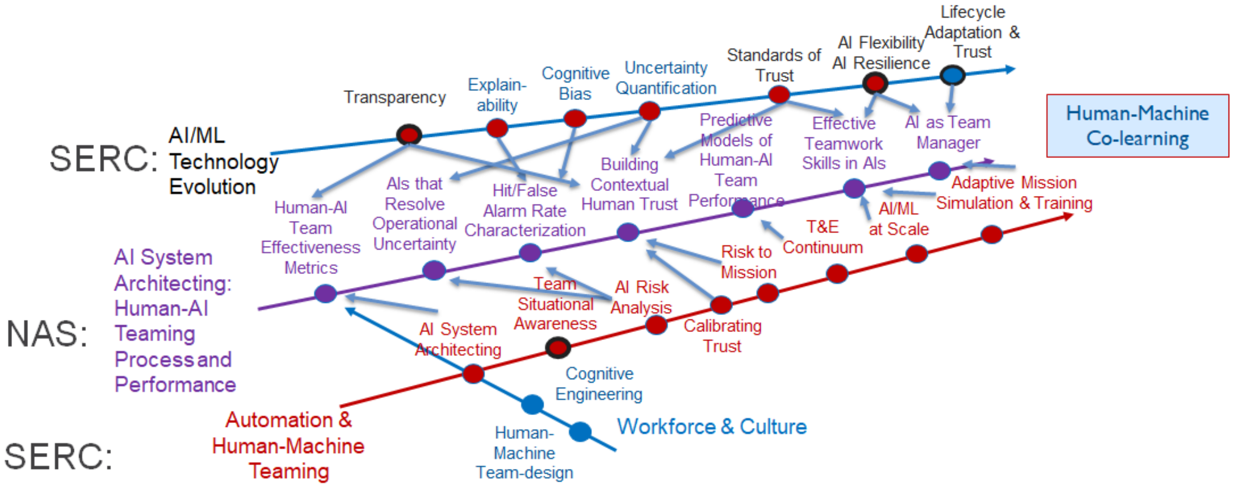

SERC:
AI/ML Technology Evolution
Transparency
Explain-ability
Cognitive Bias
Uncertainty Quantification
Standards of Trust
AI Flexibility
AI Resilience
Lifecycle Adaptation & Trust
Human-Machine Co-learning
Predictive Models of Human-AI Team Performance
Effective Teamwork Skills in AIs
AI as Team Manager
Building Contextual Human Trust
AIs that Resolve Operational Uncertainty
Hit/False Alarm Rate Characterization
Human-AI Team Effectiveness Metrics
AI System Architecting: Human-AI Teaming Process and Performance
NAS:
Adaptive Mission Simulation & Training
AI/ML at Scale
T&E Continuum
Risk to Mission
Team Situational Awareness
AI Risk Analysis
Calibrating Trust
AI System Architecting
Cognitive Engineering
Human-Machine Team-design
Workforce & Culture
SERC:
Automation & Human-Machine Teaming

Figure 3. Intersections between SERC and National Academies research.

AI Systems Architecting means the SE must shift its traditional view of architecting from a focus on structuring the system to a focus on structuring the human and system as a team. The outcome at the end of the research vectors (arrows) implies the technologies will evolve so that the AI can serve as team manager while the human serves as team leader. A primary research thrust might be articulated as building, calibrating, and setting standards for trust between Humans and AI's on the team. These are currently implemented as levels of autonomy but the SE process to design and evaluate these in real systems is much more involved than just a categorization and certification structure. This trust will differ based on criticality of systems and associated risk. Risk to Mission was added to the SERC roadmap as a result, along with Team Situational Awareness. Related to this, Transparency, Standards of Trust, AI Flexibility and Resilience, and Lifecycle Adaptation and Trust were added as research concerns in the AI/ML Technology Evolution vector. Human-machine Team Design and related Human-AI Team Testbeds were already a core focus of the SERC roadmap.

The 2023 version of the SERC roadmap represents a much stronger focus on trustworthy AI systems and factors that reflect trust. **In other words, the role of SE in the evolution of AI and Autonomy is to define and measure the qualities related to trust in today's systems.** This became one of the dominant themes for the latest 2023 SE4AI/AI4SE workshop, discussed in the next section.

The SE4AI/AI4SE Workshops

The US Army DEVCOM Armaments Center (AC) Systems Engineering Directorate (SED) and the SERC jointly sponsored the fourth AI4SE & SE4AI workshop on September 27-28, 2023, followed by a joint SERC and INCOSE virtual workshop on October 11-12, 2023. In its four years, the workshop has seen both rapid growth in presentations and attendance. From 2022 to 2023 both presentation abstracts and presented work doubled as well as attendance. Globally, over 200 attended the 2023 SERC/INCOSE virtual workshop (SERC 2020, SERC 2021, SERC 2022, SERC 2023, SERC & INCOSE 2023).

This year's workshop saw a rapid increase in applications of Large Language Models (LLMs) to partially automate traditional SE tasks like requirements development and initial model building. As these core SE tasks are heavily based on language, this is a natural entry area for AI/ML research and application. Early use of these LLM-based tools, as noted in several presentations, improve both the speed of SE tasks and the abilities of SE's with less experience to perform these tasks. It is likely that the whole "Augmented Engineering" research vector will need to be adjusted in the next several years away from technologies and tools toward exploration of how SE methods and processes can evolve to take advantage of these.

The other themes from this year's workshop echoed the vectors of the roadmap:

- Standards of trust and trustworthiness remain a dominant theme. The George Washington University (a SERC collaborator) was awarded a U.S. National Science Foundation (NSF) Research Traineeship (NRT) program entitled "Co-Design of Trustworthy AI and Future Work Systems" focused on educating researchers on future learning-based systems and methods to ensure their trustworthiness (NSF, 2021).
- The resilience of AI systems is an emerging research discussion and area, blending concepts of resilience engineering, loss-driven engineering, cyber resiliency, and digital engineering.
- Ethical considerations for AI remain an overall concern, but research is needed to bring concepts like ethics and fairness into the SE discipline as important quality attributes.
- Human-Centric approaches – the workshop highlighted the importance of considering the human element, whether in the context of users, decision-makers, or individuals interacting with AI systems. Speakers emphasized the need for understanding human cognition, decision-making processes, and the need for interpretability and explainability in AI models.

- Interdisciplinary collaboration was continuously highlighted. Within the SE discipline, the workshop highlighted collaboration between different fields such as computer science and AI/ML, systems engineering, and human systems integration. These are strongly reflected in the Workforce & Culture roadmap vector.
- Testing and experimentation, particularly in operational environments, is a lagging deployment of AI/ML and autonomous systems. The need for SE research on test, verification, and validation of these systems has become critical. There is a lack of real-world test environments, particularly to explore the relationships between data and ML.
- In the digital engineering vector, the workshop highlighted the need to establishing common infrastructure, standards, and tools for AI development and deployment across different sectors and disciplines. The workshop also called for SE research addressing the significance of data encompassing the evaluation of datasets, assurance of data quality, and the role of data in training and testing AI models.

The dialogue from the 2023 workshops underscored the imperative of adopting a holistic SE-driven approach to AI development, recognizing not only its technical dimensions but also the influence of human factors, ethical considerations, and the imperative of system interoperability. Uninterrupted testing, evaluation, and adaptation are deemed indispensable for the cultivation of robust and trustworthy AI systems, particularly within dynamic and intricate environments. Key elements in this pursuit involve red-teaming methodologies and continuous assessment. The focal point on human cognition, decision-making processes, and the psychological dynamics of AI interactions accentuates the call for a human-centered design. The seamless integration of AI into human workflows and decision processes is deemed essential for effective deployment. Additionally, ethical considerations are paramount, especially in sensitive domains such as warfare, necessitating the adoption of responsible AI practices, encompassing the development of ethical frameworks, adherence to principles, and ongoing evaluation of AI's societal impact.

Critical to engendering user trust is the model's explainability and interpretability which continues to be a technological evolution. Probably in no other domain are technology research and applications research coming together so rapidly. The workshop accentuated the significance of not only furnishing accurate outputs from AI systems but also rendering these outputs comprehensible and transparent to users. The recurring theme of ensuring the quality of training data is highlighted, emphasizing the importance of evaluating biases, reliability, and representativeness in datasets to avert adverse consequences in AI applications. Standardization of AI infrastructure, development tools, and data formats is underscored as pivotal for fostering collaboration and interoperability. These are all core SE-related challenges. The establishment of a common SE framework is deemed indispensable for the seamless integration and communication across disparate systems.

The 2023 workshop advocated for the need for calibrated metrics to gauge standards of trust in AI systems, one of the core new additions to the Roadmap. Building trust encompasses not only technical reliability but also factors such as transparency, explainability, and the capacity to adapt to dynamic conditions. Moreover, the integration of AI in military applications necessitates meticulous consideration of ethical implications, the prospect of adversarial attacks, and the development of counter-autonomy measures. Ongoing research is posited as essential to harmonize security needs with established standards.

The prominence of education and workforce development in AI-related fields is accentuated, with a call for professionals to be well-prepared to grapple with the ethical, technical, and societal challenges intrinsic to AI. Throughout all of the workshops, discussions consistently underscore the need for AI systems capable of dynamic learning and adaptation. Continuous learning, real-time adaptation, and adept decision-making in uncertain environments emerge as critical for the success of AI applications. The emphasis on active learning and adaptive systems underscores the importance of intelligently sampling observations to optimize information gain, particularly in contexts characterized by prevalent uncertainty. The practice of SE has not been fully focused on these system characteristics in the past.

In summation, all of the work presented at these workshops continues to highlight the importance in the research areas outlined in the roadmap, as well as the roadmap as a tool for discussion of these critical research and developmental needs.

Conclusion

Many disciplines, SE included, view AI and ML as new technologies. They do not clearly comprehend the breadth of impact the current growth in ML technologies on the systems we develop and the processes we use to develop them. The SERC Roadmap and the NAS report on research needs are interesting as much in their breadth of coverage as in their descriptions of individual research needs. Also, the pace of emerging AI/ML applications is fast enough that our conceptually 5-year roadmap has had to be updated every year since its original publication. For example, there were 9 presentations at the 2023 AI4SE workshop describing applications of LLMs to traditional SE functions; in 2022 there was one. The SERC will continue to publish updates to this and other SERC roadmaps to not only guide our research but also the community as a whole.

Acknowledgements

A number of people have contributed to this roadmap and the SE4AI workshops over the past 5 years. The authors would like to thank the members of the SERC Research Council who contributed, particularly Peter Beling, Valerie Sitterle, and Mark Blackburn. The authors would also like to thank the members of the INCOSE FuSE initiative and the AI Systems Working Group. As this research area continues to involve, the authors would also like to thank all the people who are exploring this active research area across the SERC and our workshop participants.

References

International Council on Systems Engineering (INCOSE) (2021). Systems Engineering Vision 2035. INCOSE. https://www.incose.org/about-systems-engineering_2/se-vision-2035_old.

Lawless, Mittu, Shortell, Sofge, and McDermott, eds. (2021). Systems Engineering and Artificial Intelligence. Springer.

National Academies of Sciences, Engineering, and Medicine. 2022. Human-AI Teaming: State-of-the-Art and Research Needs. Washington, DC: The National Academies Press. https://doi.org/10.17226/26355.

National Science Foundation (2021). Research award NRT-AI-FW-HTF: Co-Design of Trustworthy AI and Future Work Systems. https://www.nsf.gov/awardsearch/showAward?AWD_ID=2125677.

INCOSE (2020). Special Issue: AI and Systems Engineering. INSIGHT (23)1, March 2020.

Systems Engineering Research Center (SERC) (2019). Research Roadmaps 2019-2020. https://sercuarc.org/wp-content/uploads/2021/08/ROADMAPS_3.5.pdf.

SERC (2020). AI4SE and SE4AI Workshop 2020. https://sercuarc.org/event/ai4se-and-se4ai-workshop/.

SERC (2021). AI4SE and SE4AI Workshop 2021. https://sercuarc.org/event/ai4se-and-se4ai-workshop-2021/.

SERC (2022). AI4SE and SE4AI Workshop 2022. https://sercuarc.org/event/ai4se-and-se4ai-workshop-2022/.

SERC (2023). AI4SE and SE4AI Workshop 2023. https://sercuarc.org/event/ai4se-se4ai-workshop-2023/.

SERC and INCOSE (2023). AI4SE & SE4AI Virtual Workshop 2023. https://sercuarc.org/event/serc-incose-ai4se-se4ai-workshop-2023/

Biographies

Tom McDermott. Tom McDermott is the Chief Technology Officer of the Systems Engineering Research Center (SERC) and a faculty member in the School of Systems and Enterprises at Stevens Institute of Technology in Hoboken, NJ. With the SERC he develops new research strategies and is leading research on digital transformation, education, security, and artificial intelligence applications. He previously held roles as Faculty and Director of Research at Georgia Tech Research Institute and Director and Integrated Product Team Manager at Lockheed Martin. Mr. McDermott teaches system architecture, systems and critical thinking, and engineering leadership. He provides executive level consulting as a systems engineering and organizational strategy expert. He is a fellow of the International Council on Systems Engineering (INCOSE) and recently completed 3 years as INCOSE Director of Strategic Integration.

Kara Pepe. Kara M. Pepe is the Director of Operations at the Systems Engineering Research Center (SERC). Prior to joining SERC, Kara was the Director of Industry and Government Relations at the Center for Complex Systems and Enterprises, working with the various government, private sector, and non-profit organizations that engaged and funded research initiatives for the center. She received her M.E. in Systems Engineering and B.E. in Engineering Management from Stevens and is currently pursuing a PhD. Her research focuses on digital engineering in general with an emphasis on workforce development. Kara is a member of INCOSE, NDIA, SWE, and ASEM.

Megan Clifford. Megan M. Clifford is a Research Associate and Engineer at Stevens Institute of Technology. She works on various research projects with a specific interest in systems assurance, cyberphysical systems, and programs with national and global significance. She previously worked on the leadership team as the Chief of Staff and Program Operations for the Systems Engineering Research Center (SERC), was the Director of Industry and Government Relations to the Center for Complex Systems and Enterprises (CCSE), and held several different positions, including Systems Engineer, at Mosto Technologies while working on the New York City steam distribution system.