The government seeks individual input; attendees/participants may provide individual advice only.

Middleware and Grid Interagency Coordination (MAGIC) and AI R&D IWG Joint Meeting Minutes
September 6, 2023
Virtual Meeting

Participants

Alex May (ORNL)	Kimberly Sablon (DOD/OSD)
Alex Phounsavath (DHS S&T)	Kyle Fox (NIJ)
Alison Derbenwick Miller (Oracle)	Laura C (NSA)
Allison Dennis (NIH)	Devon Diaz (Army AI2C)
Andrew Merluzzi (USAID)	Mallory Hinks (NCO)
Ann Stapleton (USDA OCS)	Marcy Collinson (Oracle)
Anthony Cruz (DOD/OSD)	Mark Walderhaug (FDA)
Arjun Shankar (ORNL)	Marlon Pierce (NSF) (Marlon Pierce)
Brian Henz (DHS S&T)	Matthew Abernathy (DOD/MDA)
Christyann Pulliam (USPTO)	Michael Littman (NSF)
Dan Cosley (NSF)	Mubassira Khan (DOT)
Dave Barnes (Army AI2C)	Neeti Pokhriyal (NSF)
David Etim (DOE/NNSA)	Nelson Colon Vargas (FDA)
David Kuehn (DOT/FHWA)	Nikunj Oza (NASA)
Eric Lancon (BNL)	Robin Brown (DIA/CTO)
Faisal DSouza (NCO)	Rajeev Agrawal (DASA-P)
Hsin Fang (NIST)	Rajendra Raj (Rochester Institute of Technology)
Feiyi Wang (ORNL)	Ralph Wachter (NSF)
Gabriel Martinez (DHS CISA)	Ram D. Sriram (NIST)
Gail-Joon Ahn (ASU)	Ravi Madduri (ANL)
Garfield Jones (DHS CISA)	Robert McFarland (OUSD-R&E)
H Birali Runesha (University of Chicago)	Roger Miller (NIDCD)
Hal Finkel (DOE SC)	Shantenu Jha (BNL, Rutgers)
Hubertus van Dam (BNL)	Sharon Geva (NSF)
Ian Karlin (NVIDIA)	Shawn Forrest (FDA)
Iwona Weidlich (FDA HHS)	Srini Mandalapu (FAA)
Jack Wells (NVIDIA)	Stephanie Garcia (HHS/ONC)
Jay Vietas (NIOSH)	Steven Lee (DOE/SC)
Jeff Larkin (NVIDIA)	Craig Thor (DOT/FHWA)
Jerry Ma (USPTO/ATR)	Thuc Hoang (DOE/NNSA)
Jessica Li	Tom O'Neill (DOE)
John Garofolo (NIST)	Tomas Drgon (FDA)
Kamie Roberts (NITRD)	Tony Cruz (OUSD R&E)
Keith Beattie (LBL)	

Introductions: This joint meeting was chaired by Jay Park (NSF) and Hal Finkel (DOE SC), MAGIC co-chairs, and Steven Lee (DOE/SC) and Michael Littman (NSF), co-chairs for the AI R&D IWG.

Implementing Responsible AI across the DOD Research Enterprise

Kimberly Sablon (DOD OUSD R&E)

Dr. Sablon discussed the evolution of AI policy and strategy, the AI ethical principles outlined in the 202 AI Strategy and Implementation Plan, implementation tenets, and R&E's efforts in implementing responsible AI. She highlighted the role of research and equity and ethics, as well as the history of the DOD AI strategy. She emphasized the importance of warfighter trust and verification/validation (V&V) in AI systems.

She discussed the purpose and objectives of the Center for Calibrated Trust Measurement and Evaluation (CAIT), including its focus on operationalizing responsible AI and advancing the science of test evaluation V&V of AI and autonomous systems. She also highlighted the importance of ethics, security, and the warfighter-in-the-loop design, as well as collaboration with academic institutions and international partnerships to develop a curriculum for responsible AI test and evaluation.

Dr. Sablon also talked about the implementation of a testbed for continuous adversarial testing and red teaming in AI systems. She mentioned the development of repositories, performance risk metrics, and addressing data integration issues. She also emphasized the importance of ethics and value alignment in the design stage of AI systems.

She highlighted the challenges of implementing ethics and laws in AI and robotic systems at the design stage. She mentioned the insufficiency of both top-down and bottom-up approaches and emphasized the need for a value-driven logic that integrates ethics, moral principles, and decision-making into AI systems. She also mentioned ongoing discussions with small businesses and larger companies on integrating value-driven logic into AI system design.

Dr. Sablon discussed the design of AI agents that can serve as artificial moral agents, making ethical moral decisions and leading to ethical moral behaviors. She mentioned the estimation of values, assessment of reliability, calculation of rewarding/punishing effects, facilitation of learning, and inclusion of behavioral constraints from rules, laws, and ethics. She also emphasized the importance of data in informing the way forward.

She talked about the importance of involving warfighters in the design and evaluation of AI systems through tabletop exercises and qualitative data assessments. She emphasized the need for well-informed and robust simulations to quantify metrics of trust and understand factors that drive warfighters' trust levels. CAIT is being developed with MIT Lincoln Lab and Carnegie Mellon to shape this process. She also mentioned the role of research in equity and ethics and highlighted the efforts of organizations like AFOSR and ONR in investigating human perceptions, beliefs, behaviors, and computational architectures for robots that will contribute to CAIT.

Dr. Sablon also mentioned various ethical features in AI research, including explainable AI and assured autonomy. She talked about the DARPA ERSA program focused on developing a legal, moral, and ethical framework for autonomous sensing. She also highlighted the Missile Defense Agency (MDA) AGES program's work on human-machine teaming and ethical guidelines set by the Defense Innovation Unit (DIU) for working with the commercial sector. She emphasized the commitment to ethical and responsible use of AI while acknowledging the need to further define what responsible AI means.

Questions

- Stephanie Garcia from HHS, ONC asked "How is this different than a cost benefit analysis? Is the moral/ethical decision-making still left up to a human? And this is providing supporting data to help a human make a decision?"
    - Dr. Sablon: "Decision support in many instances, it's not the autonomous system making the decision. That's why there's a human in the loop. But again, when I talk about the value component, to me it comes down to having a process that can compute value, determine its importance of that value, assess reliability, and generate reward or punishment to that reward based on desirable behavior or punish on desirable behavior and so forth. On the autonomy side, when we're talking human in the loop, sure, you know, you're doing a quick analysis and you're providing, you know, COAs that the human can take into consideration. But we want to make sure, too, that if there's an autonomous system that's scouting out there, we're making where the system can operate in a very responsible manner because now we have the right volumes and values in place and the ability to compute those values and assess reliabilities."
- John Garfolo from NIST asked "How would you test moral agents?"
    - Dr. Sablon: "That's a good question. I don't have an answer for that yet. We're still working through, you know, again, what those moral principles are and what value-driven logic looks like. How do you even encode that, you know, put that into code and then start coming up with clear metrics to test them? In large part, that's what's going to be figured out under CAIT."
- Mubasira Khan asked "What data will be used to train the moral agents? How would you ensure the data is not biased?"
    - Dr. Sablon: "Again, it's going to be operational data combined with some synthetic data. That's why again the human in the loop, we're doing a lot of the data collection, like I talked about when we're putting all of this in simulation under CAIT. Much of the data collection is going to come from tabletop exercises. We're pulling operational data to, you know, that's in theatre. The best operational data that we could utilize would be data that actually collects from war. We've only had so much war, so again it's going to be a trusted data collection. It's what we're going to do in both virtual and live fire environments, and we're going to do that utilizing warfighter touchpoints. Again, from tabletop exercises we want to make sure we're designing some of the experiments, you know, have some analysis of what human trust surveys and interviews qualitatively looks like."

- Rajiv Agrawal asked "If allowed to share, who is the lead PI of this project?"
    - Dr. Sablon said CAIT is going to be led out of CMU.

NITRD AI R&D Strategic Plan – 2023 Update

Steven Lee (DOE/SC), Michael Littman (NSF), Craig Schlenoff (NIST), AI R&D IWG Co-chairs

Dr. Lee provided an overview of the 2023 update to the NITRD AI R&D Strategic Plan. He discussed the purpose of the plan, the previous 2019 update, and the gathering of information for the 2023 update. He also mentioned the involvement of various agencies, the release date of the report, and provided a visual representation of the collected data from the public RFI. Additionally, he introduced the leads and contributors involved in the writing of the strategies and presented a rough partitioning of the R&D foundations.

Dr. Lee provided an overview of each of the nine strategies outlined in the NITRD AI R&D Strategic Plan update.

- Strategy 1: Make long-term investments in AI research.
    - This gets at prioritizing investments in the next generation of AI to drive responsible innovation. This includes foundational capabilities and federated machine learning approaches. He also highlighted the need for research on scalable artificial general intelligence systems.
- Strategy 2: Develop effective methods for human-AI collaboration.
    - This is for increasing our understanding of the attributes of successful human-AI teams and mitigating the risk of human misuse.
- Strategy 3: Ethical, legal, and societal implications of AI.
    - Dr. Lee emphasized the importance of developing approaches to mitigate risks and promote equity through interdisciplinary research.
- Strategy 4: Ensure the safety and security of AI systems.
    - Advance knowledge of how to design AI systems that are trustworthy, reliable, dependable, and safe. This includes research to advance the ability to test, validate, and verify the functionality and accuracy of AI systems, and to secure AI systems from cybersecurity and data vulnerabilities.
- Strategy 5: Develop shared public datasets and environments for AI training and testing.
    - Develop and enable access to high-quality data sets and environments, as well as to testing and training resources, a broader and more diverse community engaging with the test data and tools for conducting AI research, increases the potential for more innovative and equitable results.
- Strategy 6: Measure and evaluate AI systems through standards and benchmarks.
    - Develop a broad spectrum of evaluated techniques of AI, including technical standards and benchmarks that are informed by the administration's blueprint for AI Bill of Rights, and also informed by the AI risk management framework coming from NIST.
- Strategy 7: Better understand the national AI R&D workforce needs.

    - Improve opportunities for R&D workforce development to strategically foster an AI-ready workforce in America. This includes R&D to improve understanding of the limits and possibilities of AI and AI-related work and the education and fluency needed to effectively interact with AI systems.
- Strategy 8: Expand public-private partnerships to accelerate advances in AI.
    - Promote opportunities for sustained investment, responsible AI, R&D, and for transitioning advances into practical capabilities in collaboration with academia, industry, international partners, and other non-federal entities.
- Strategy 9: Establish a principled and coordinated approach to international collaboration in AI research
    - Prioritize AI R&D with our international partners to address global challenges such as environmental sustainability, advances into practical capabilities in collaboration with the academia and industry.

Dr. Lee discussed the new metrics included in the 2023 update of the NITRD AI R&D Strategic Plan. These metrics aim to evaluate the implementation of the National AI Initiative Act and Strategic Plan by federal agencies. The metrics include tracking investments in AI R&D, education, and workforce development, as well as monitoring multi-agency programs, diversity of active users in data sets, and federally supported AI test beds.

<u>Questions</u>

- Marlon Pierce (NSF): I'm glad to see Strategy 9. How do you undertake this while also addressing issues of national security and primacy, particularly with nations with which we do not have good relations?
    - Michael Littman said a lot of what is happening is very pairwise. There's particular arrangements that are being made with particular countries. Those arrangements typically take into consideration exactly the sort of sensitivities Marlon is alluding to.
    - Steven Lee said that he concurred. There are targeted collaborations with the EU or UK and other allies, rather than a free for all approach.
- John Garofolo (NIST) asked is there a difference between core AI R&D priorities and R&D to support applied national needs. Those don't seem to have been separated.
    - Michael Littman said that he thinks of those as being separate, but also quite interrelated.
    - Steven Lee said that the plan deals with the cross-cutting needs and the foundations for that, and some will be relevant to both. But some may be more specific to applied national needs. He said that they did take their foundational research approach. What are the cross-cutting areas and or the areas of long-term research investments. So they're not necessarily separate.
- Mubassira Khan asked do we have a good understanding about the current state of our organizations' infrastructure needed for AI model training and testing. Does Strategy 5 include that topic?

MAGIC Team Overview
Hal Finkel (DOE/SC) MAGIC Co-chair

Dr. Finkel explained the purpose and history of MAGIC (Middleware and Grid Interagency Coordination Team). He discussed the focus of MAGIC on middleware and grid computing resources and highlighted the importance of coordination and standardization. He also mentioned the previous co-chair, Jay Park, the structure and topics of the meetings, and the tradition of an annual in-person meeting at SC with a focus on cybersecurity. Additionally, he mentioned recent discussions on multitenancy at different levels and presentations from national laboratories.

Dr. Finkel provided an overview of the recent presentations and discussions on multi-tenancy and workflows. He mentioned presentations from Oracle and Microsoft on multitenancy in cloud infrastructure, as well as presentations on resiliency and multitenancy in high-performance computing. He also mentioned presentations on AI workflows and managing resources for both batch and urgent workflows. He also discussed future architectures and their impact on grid and distributed computing, data-centric computing resources, and composability across infrastructures.

Dr. Finkel talked about upcoming discussions and activities in MAGIC, including an upcoming discussion on the NSF AI Institute for Network Systems. He mentioned the possibility of inviting a speaker from OSTP and conducting more joint interagency working group meetings. He also highlighted the open nature of the group, welcoming feedback from participants and mentioning other potential topics of interest for future discussions.

Questions:

- Steven Lee asked if MAGIC had a strategic plan or priorities.
    - Mallory said MAGIC is situated under the Large-Scale Networking IWG. MAGIC itself doesn't have strategic priorities. MAGIC has a rough plan of what topics we want to cover, and we report that information to LSN during the APM.
    - Hal said MAGIC is a public group and we do not make recommendations or strategic plans for ourself or for any other agency, but we do serve as a forum for open discussion.
- Matthew Abernathy asked if anybody from the DOD's, HPCMP participate in MAGIC?
    - Mallory and Hal were not sure.
    - Matthew said he could reach out to tell people about MAGIC.

Open Discussion:
The participants discuss various topics related to AI research and funding. Hal mentioned the efforts of DOD, NIST, and DOE in the field of AI. He also mentioned voluntary commitments on AI announced by the White House. Hal asked if anyone had received a research paper or funding proposal written by an AI system. Matthew recalled a related example from USPTO. Tomas shared an example from FDA.

The participants discussed the presence of AI-generated applications in various submissions. Tomas mentioned a spike in comments related to AI and suggested grouping them. Iwona noted that AI has been used in submissions related to drugs and medical devices, but the information is not publicly available.

Christyann from USPTO mentioned that they have not publicly acknowledged the presence of applications generated by AI. She discussed a recent case in which an AI system was listed as the inventor but the Supreme Court denied the patent. Alison mentioned that ACM has a policy against AI authors in their digital library. Hal and Alison discussed the challenges of controlling AI and ensuring it stays within intended boundaries.

Matthew asked the MAGIC group about their work on preserving security and separation between systems in cloud deployments. Hal and Mallory discussed previous presentations on federation and isolation between systems for security purposes. Matthew mentioned the work of CDAO on federating different networks across the DOD.

Next Meeting November 14, 2023 at SC23