



ARL-SR-0487 • DEC 2023



2023 Hybrid Thinking Workshop: Insights on Future Hybrid Human–Machine Intelligence and Army Command and Control

by Steven M Thurman, Javier O Garcia, Luis Jimenez,
Christopher Kelshaw, and Tom Christensen

DISTRIBUTION STATEMENT A. Approved for public release: distribution unlimited.

NOTICES

Disclaimers

The findings in this report are not to be construed as an official Department of the Army position unless so designated by other authorized documents.

Citation of manufacturer's or trade names does not constitute an official endorsement or approval of the use thereof.

Destroy this report when it is no longer needed. Do not return it to the originator.



2023 Hybrid Thinking Workshop: Insights on Future Hybrid Human–Machine Intelligence and Army Command and Control

Steven M Thurman and Javier O Garcia
DEVCOM Army Research Laboratory

Luis Jimenez
Oak Ridge Associated Universities

Christopher Kelshaw and Tom Christensen
Mission Command Battle Lab (MCBL)

REPORT DOCUMENTATION PAGE

1. REPORT DATE		2. REPORT TYPE		3. DATES COVERED	
December 2023		Special Report		START DATE	END DATE
				12 June 2023	14 June 2023
4. TITLE AND SUBTITLE					
2023 Hybrid Thinking Workshop: Insights on Future Hybrid Human–Machine Intelligence and Army Command and Control					
5a. CONTRACT NUMBER		5b. GRANT NUMBER		5c. PROGRAM ELEMENT NUMBER	
5d. PROJECT NUMBER		5e. TASK NUMBER		5f. WORK UNIT NUMBER	
6. AUTHOR(S)					
Steven M Thurman, Javier O Garcia, Luis Jimenez, Christopher Kelshaw, and Tom Christensen					
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)				8. PERFORMING ORGANIZATION REPORT NUMBER	
DEVCOM Army Research Laboratory ATTN: FCDD-RLA-FB Adelphi, MD 20783				ARL-SR-0487	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSOR/MONITOR'S ACRONYM(S)	11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT					
DISTRIBUTION STATEMENT A. Approved for public release: distribution unlimited.					
13. SUPPLEMENTARY NOTES					
ORCID IDs: Steven M Thurman, 0000-0003-3962-8447; Javier O Garcia, 0000-0001-5459-8974					
14. ABSTRACT					
<p>The US Army Combat Capabilities Development Command Capabilities Development Command (DEVCOM) Army Research Laboratory (ARL) is developing an anti-disciplinary 6.1 basic research program on hybrid human–machine intelligence that seeks new opportunities and methods to synergize human intelligence and machine intelligence to enable advanced capabilities for military applications. To this end, ARL researchers organized a 3-day Hybrid Thinking Workshop that was hosted at Stanford University and brought together top academic researchers from a myriad of academic disciplines and Army subject matter experts in command and control with ARL researchers to share information and ideate on future applications of hybrid human–machine intelligent systems to specific challenges in command and control. This report summarizes the purpose, aims, and design of the Hybrid Thinking Workshop, and it summarizes the presentations provided by invited academic researchers and concept developers from Mission Command Battle Lab. Finally, the report discusses prominent concepts and ideas that emerged during the group discussions and breakout sessions associating hybrid intelligent machine-aided technologies to various aspects of the Army’s command and control system.</p>					
15. SUBJECT TERMS					
Humans in Complex Systems, artificial intelligence, command and control, hybrid intelligence, neuroscience, Human–Machine Teaming, Workshop					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT		18. NUMBER OF PAGES
a. REPORT	b. ABSTRACT	c. THIS PAGE	UU		44
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED			
19a. NAME OF RESPONSIBLE PERSON				19b. PHONE NUMBER (Include area code)	
Steven M Thurman				(301) 488-5390	

Contents

List of Figures	v
Acknowledgments	vi
Executive Summary	vii
1. Introduction: Purpose and Aims of the Hybrid Thinking Workshop	1
1.1 Design of the Hybrid Thinking Workshop	2
1.2 Expected Outcomes of the Workshop	5
2. Summary of Presentations	6
2.1 Academic Talks	6
2.1.1 Dr Christopher Re	6
2.1.2 Dr Boyuan Chen	7
2.1.3 Dr Ramesh Srinivasan	8
2.1.4 Dr Janet Rafner	9
2.1.5 Dr Gillian Smith	9
2.1.6 Dr Manish Saggar	10
2.1.7 Dr Stephen Fiore	11
2.1.8 Dr Lena Smirnova	12
2.1.9 Dr Michael Beyeler	13
2.1.10 Dr Sarah Kitchen	14
2.1.11 Dr Ranjay Krishna	15
2.2 Army SME Talks: Mission Command Battle Lab	16
2.2.1 Day 1	17
2.2.2 Day 2	21
2.2.3 Day 3	24
3. Conclusion	26
3.1 Future US Army C2 Should Require Relatively Fewer People and Organizations	27
3.2 Future US Army C2 Should Allow Commanders to Make Faster, Less Error-Prone, and More Insightful Decisions Than the Enemy	28

3.3	Future US Army C2 Should Function Well, Even When Command Post Nodes Are Dispersed and, at Times, Moving	29
3.4	Priority Research Questions Going Forward	30
	List of Symbols, Abbreviations, and Acronyms	33
	Distribution List	35

List of Figures

Fig. 1	Image of the Hybrid Thinking Workshop agenda for Day 1	3
Fig. 2	Image of the Hybrid Thinking Workshop agenda for Day 2	4
Fig. 3	Image of the Hybrid Thinking Workshop agenda for Day 3	5

Acknowledgments

This research is aligned with the scientific aims of the Humans in Complex Systems Division of the US Army Combat Capabilities Development Command (DEVCOM) Army Research Laboratory (ARL). We want to acknowledge the intellectual contribution of ARL scientists and engineers and their strong influence on the scientific development of the Hybrid Human-Technology Intelligence program. We also acknowledge the significant contribution of the professional staff and leadership at Mission Command Battle Lab, Fort Leavenworth, Kansas, to the workshop described in this report. Finally, we thank and acknowledge the invited academic scientists that presented their research, in-person and remotely, and all workshop attendees that contributed to discussions at the workshop.

Executive Summary

The US Army Combat Capabilities Development Command (DEVCOM) Army Research Laboratory's (ARL's) Humans in Complex Systems Division is developing a basic research program in the area of hybrid human-machine intelligence to study new opportunities and methods to synergize human and machine intelligence to enable future capabilities beyond what a human or machine could provide alone. On June 12–14, 2023, ARL organized a 3-day workshop hosted at Stanford University to bring together academic experts in a variety of domains related to hybrid thinking (neuroscience, artificial intelligence [AI], human-agent teaming, cognitive science, applied math), ARL scientists, and US Army subject matter experts from Mission Command Battle Lab (MCBL) on the future of command and control (C2). The schedule of the workshop was designed to enable information sharing and discussion of various research topics, as well as an overview of the Army C2 system. During afternoon breakout sessions, all attendees split into smaller groups to collaboratively ideate on applications of AI, automation, and hybrid technologies to fulfill specific future capability objectives identified by MCBL. This report provides an overview of the purpose and aims of the workshop, a description of the invited attendees and the work that was presented, and a conclusion that highlights relevant future technology areas and critical research questions.

1. Introduction: Purpose and Aims of the Hybrid Thinking Workshop

The US Army Combat Capabilities Development Command (DEVCOM) Army Research Laboratory's (ARL's) Humans in Complex Systems Division launched the Hybrid Human-Technology Intelligence (HHTI) program in FY23 to study new methods to combine artificial intelligence and machine learning (AI/ML) technologies with unique aspects of human intelligence to develop new systems for human-agent teaming in future Army contexts. Future operating environments will be more fast-paced and complex than humans can handle, while simultaneously requiring more adaptation, creativity, and leadership than machines can provide. We must consider hybridized human-technology systems built on inputs from both human and machine intelligence that operate symbiotically to enable new battlefield capabilities to keep up the pace and effectively accomplish mission goals.

Our core thesis is that human intelligence effectively partnered with machine intelligence will win the intelligence race by creating adaptive, moral human-technology teams that outpace and outlast other forms of technology. While human thought can be overwhelmed with too much information arriving too quickly, human thinking is also uniquely creative and adaptable in producing sound approximations in novel and open-ended contexts. On the other hand, machines can handle high volumes of data but can fail catastrophically when exposed to novel inputs. Basic research is needed to understand how to maximize the strengths, while minimizing the weaknesses, of hybrid human-technology systems to generate novel capabilities for the Army.

The general scientific approach at the outset of the HHTI program is to aggregate information from a variety of academic labs that have led recent advancements in these areas, as well as Army leaders and subject matter experts (SMEs), to characterize the limitations of both human and machine intelligence, understand preliminary capabilities of hybrid intelligence systems, and establish an Army-relevant experimental platform. To this end, the "Hybrid Thinking" Workshop was organized to bring together academic experts, Army SMEs, and ARL researchers for a 3-day event to interact in person, share information, and ideate on future hybrid technologies and applications in the area of command and control (C2). The design and organization of this workshop was a core deliverable of the HHTI program for FY23 to help establish a footing for the program to grow into the future.

1.1 Design of the Hybrid Thinking Workshop

On June 12–14, 2023, the Hybrid Thinking Workshop was held at the Stanford University Faculty Club. This 3-day event featured academic presentations and group discussions with interactive breakout sessions. The purpose of this event was to bring thought leaders across academia, industry, and the DOD to merge novel advances in AI and ML with advances in neuroscience, behavioral science, and related areas to advance our understanding of synergistic, hybrid human–machine intelligence.

Attendees actively participated with Army scientists and concept developers to enable foundational research programs in this area for future hybrid technologies. Due to the interdisciplinary, or even anti-disciplinary, nature of the HHTI research program, the workshop brought together a diverse range of academic scientists with expertise in fields including computer science, ML, cognitive science, neuroscience, human–agent teaming, applied mathematics, and hybrid intelligence.

The schedule of the workshop was designed to encourage information sharing, collaborative exchange, innovative thinking, and collaborative ideation on future Army applications for Hybrid Thinking research products. As shown in Figs. 1–3, each day of the workshop was organized as follows:

- 1) A morning session that included 30-min presentations from three to four invited academic researchers. The morning format allowed ARL scientists on the East Coast that could not attend in person to view the talks, ask questions, and participate remotely.
- 2) An afternoon session that included a 60–90 min presentation from Army C2 SMEs (Christopher Kelshaw, Tom Christensen) from Mission Command Battle Lab (MCBL), Fort Leavenworth, Kansas. The afternoon session was also live-streamed to ARL researchers for remote participation.
- 3) An interactive breakout session for 120 min to ideate in small groups on hybrid intelligent applications to Army C2 processes and challenges. The breakout sessions were exclusive for in-person attendees.

Hybrid Human-Technology Intelligence Research Symposium				
Day One: 12 June 2023, Stanford Faculty Club, 439 Lagunita Dr., Stanford, CA 94305 (all times are PST)				
Start	Time	End	Title	Briefers
8:30	0:30	9:00	Arrival at Stanford Faculty Club	
9:00	1:00	10:00	Open Floor Continental Breakfast	
Informational Presentations				
10:00	0:05	10:05	General Introduction	Steven Thurman, DEVCOM ARL
10:05	0:10	10:15	DEVCOM ARL HCxS Overview	Kaleb McDowell, DEVCOM ARL
10:15	0:15	10:30	HHTI Program Overview	Javier Garcia, DEVCOM ARL
Decision Making, Crowdsourcing, Human-machine Teaming				
10:30	0:30	11:00	Keynote Opening Talk: Foundational Models	Chris Re, Stanford University
11:00	0:30	11:30	General Purpose Human-AI Teaming Framework	Boyuan Chen, Duke University
11:30	0:30	12:00	Neurocognitive Models for Hybrid Thinking	Ramesh Srinivasan, University of California Irvine
12:00	1:00	13:00	Catered Lunch at Stanford Faculty Club	
Breakout Activities				
13:00	0:30	13:30	Primer on Command and Control (C2)	Mission Command Battle Labs
13:30	1:00	14:30	Group Breakout Session	
14:30	0:20	14:50	Coffee/Snack Break	
14:50	0:50	15:40	Group Breakout Session II	
15:40	0:20	16:00	General Group Discussion and Wrap Up	
16:00			Adjourn	

Fig. 1 Image of the Hybrid Thinking Workshop agenda for Day 1

Hybrid Human-Technology Intelligence Research Symposium				
Day Two: 13 June 2023, Stanford Faculty Club, 439 Lagunita Dr., Standord, CA 94305 (all times are PST)				
Start	Time	End	Title	Briefers
8:30	0:30	9:00	Arrival at Stanford Faculty Club	
9:00	0:30	9:30	Open Floor Continental Breakfast	
Creativity, Human AI Teaming				
9:30	0:30	10:00	Hybrid Intelligence and Human-AI Co-Creativity	Janet Rafner, Aarhus University
10:00	0:30	10:30	Understanding the Social Impact of Generative AI via Critical Design	Gillian Smith, Worcester Polytechnic University
10:30	0:30	11:00	ideAlte: A Synergistic Human-Machine Framework for Ideation and Prototyping	Manish Saggarr, Stanford University
11:00	0:30	11:30	Conceptualizing Bio-Technical Intelligence for Human-Machine Systems	Stephen Fiore, University of Central Florida
11:30	1:00	12:30	Catered Lunch at Stanford Faculty Club	
Breakout Activities				
12:30	0:30	13:00	C2 Processes	Mission Command Battle Labs
13:00	1:00	14:00	Group Breakout Session	
14:00	0:20	14:20	Coffee/Snack Break	
14:20	1:00	15:20	Group Breakout Session II	
15:20	0:40	16:00	General Group Discussion and Wrap Up	
16:00			Adjourn	

Fig. 2 Image of the Hybrid Thinking Workshop agenda for Day 2

Hybrid Human-Technology Intelligence Research Symposium				
Day Three: 14 June 2023, Stanford Faculty Club (all times are PST)				
Start	Time	End	Title	Briefers
8:30	0:30	9:00	Arrival at Stanford Faculty Club	
9:00	0:30	9:30	Open Floor Continental Breakfast	
Neurotechnology, Modelling, Applied Mathematics				
9:30	0:30	10:00	Organoid Intelligence: Can We Replace Animal Models of Learning and Memory with Brain Organoids	Lena Smirnova, Johns Hopkins University
10:00	0:30	10:30	Human-in-the-Loop Machine Learning for Neurotechnologies	Michael Beyeler, University of California, Santa Barbara
10:30	0:30	11:00	Dynamical Systems and Interacting Population Models on Graphs	Sarah Kitchen, Michigan Tech Research Institute
11:00	0:30	11:30	Machine learning from human interactions	Ranjay Krishna, University of Washington
11:30	1:00	12:30	Catered Lunch at Stanford Faculty Club	
Breakout Activities				
12:30	0:30	13:00	Pattern of Work for a Staff Officer	Mission Command Battle Labs
13:00	1:00	14:00	Group Breakout Session	
14:00	0:20	14:20	Coffee/Snack Break	
14:20	0:40	15:00	Group Breakout Session II	
15:00	0:30	15:30	General Group Discussion and Wrap Up	
15:30	0:30	16:00	Final Research Symposium Wrap Up	
16:00			Adjourn	

Fig. 3 Image of the Hybrid Thinking Workshop agenda for Day 3

1.2 Expected Outcomes of the Workshop

The primary expected outcomes of the workshop were 3-fold. First, to gather an anti-disciplinary group of academic experts to consult and collaborate with ARL scientists to identify critical scientific questions and outline a research program that can set us on a path to progress in a complex problem space at the boundary of our current paradigm of interaction with technology. As we build an Army-centered research ecosystem in this emerging research area, it will be critical to establish relationships with top-tier academic labs to develop collaborative research efforts and support recruitment to enrich the future ARL workforce in this area.

The second aim was to expose the group of academic scientists, as well as ARL scientists, to information about significant current and future Army challenges in the domain of C2. The purpose of the interaction with future Army C2 SMEs was to inspire and ground our group discussions of future technology in the application space of Army commanders and their staff to support the decision-making process on the battlefield. Academic researchers rarely get the direct insight and access to Army SMEs that was provided by this workshop. The information provided and following discussions may help inform and shape their future research agendas and increase their interest and ability to work directly with ARL researchers or other DOD labs in research programs like HHTI.

The third aim was to generate a technical report that summarizes the activities of the workshop and presents the core ideas from breakout sessions to guide the direction and scope of future research efforts within ARL and the DOD with regards to human–intelligent machine teaming in support of future C2. This report is the third deliverable of the workshop.

2. Summary of Presentations

The following section introduces the biographical sketches of each invited academic presenter and provides an overview of their research presentations.

2.1 Academic Talks

2.1.1 Dr Christopher Re

Christopher (Chris) Re is an associate professor in the Department of Computer Science at Stanford University. He is in the Stanford AI Lab and is affiliated with the Statistical Machine Learning Group and the Center for Research on Foundation Models. His recent work is to understand how software and hardware systems will change as a result of ML along with a continuing, petulant drive to work on math problems. Research from his group has been incorporated into scientific and humanitarian efforts, such as the fight against human trafficking, along with products from technology and enterprise companies including Apple, Google, YouTube, OpenAI, and more.

Dr Re’s talk discussed how large-scale organizational functions are being replaced by AI, effectively changing the nature of leadership and responsibilities in business. For instance, large companies like Apple have made efficient use of a self-supervised model that can match places to a database in over 40 languages without direct translations. These models, trained on public data like Wikipedia, outperformed Apple’s previous system, and required much less effort.

Organizations are now moving to a “copilot era” where people augment their work with AI assistants. Tasks will be replaced and workers reassigned to meet new demands or be retrained. However, AI-generated content can contain erroneous facts, despite being better than humans in testing. There remains a significant gap between intended and actual deployment of AI systems, and the need for validation and basic research. Another limitation of traditional computer science and ML approaches is that they narrow down problems to specific areas and train models to predict specific outcomes. Dr Re’s research group is currently exploring the use of generalist models that can handle a wide range of tasks and data.

2.1.2 Dr Boyuan Chen

Boyuan Chen joined Duke University in 2022 as an assistant professor in the Department of Mechanical Engineering and Materials Science, the Department of Electrical and Computer Engineering, and Department of Computer Science. Dr Chen obtained his PhD and MS in Computer Science at Columbia University with Professor Hod Lipson. He obtained his BS in Electrical Engineering and Biomedical Engineering from Jilin University. His research focuses on robotics, computer vision, ML and dynamical system modeling. He is interested in developing “generalist robots” that learn, act, and improve by perceiving and interacting with the complex and dynamic world. His research often relies on the natural and unlabeled sensory inputs from multiple modalities. Ultimately, he hopes that robots and machines can equip with high-level cognitive skills to assist people and unleash human creativity.

Dr Chen’s talk focused on developing AI that can understand different roles, communicate effectively with humans, and collaborate as part of a team. Recently, his lab built a platform called “CREW” to study human–AI collaboration at a large scale. CREW allows researchers from different backgrounds to study human–AI collaboration using a common platform. Dr Chen demonstrated how CREW is currently being used to run a hide and seek game with dozens of agents. Users can join or host sessions, view or provide feedback to agents, directly control agents, select roles, and view different perspectives. Parallel message sessions can be embedded in environments, and agent behaviors can be customized to embrace unique roles. The CREW platform overcomes limitations of current AI systems by allowing researchers to track and record multi-human and multi-agent behavior in large-scale, realistic games. Moreover, the multi-player environments allow for massive human feedback to train AI agents governed by various learning algorithms. This engaging, flexible, and extensible platform allows a diverse set of researchers to study AI with theory of mind, human decision making, collaboration, hybrid thinking, and team dynamics.

2.1.3 Dr Ramesh Srinivasan

Dr Srinivasan's PhD training in BME at Tulane University was on the development of theoretical models of electroencephalography (EEG). He developed cognitive science and neuroscience expertise through postdoctoral training at the University of Oregon and as a Fellow in Theoretical Neurobiology at the Neurosciences Institute. He joined the faculty in Cognitive Sciences and Biomedical Engineering in 2000 where he has continued his combined theoretical and experimental research. He was Chair of Cognitive Sciences from 2012 to 2022. The primary focus of Srinivasan's research is on developing methodologies to extract network structure from EEG, magnetoencephalography (MEG), and functional magnetic resonance imaging (fMRI) signals and relate it to cognitive functions. This foundational signal processing research has had a broad impact in clinical and cognitive neuroscience research. Most pertinent to the theme of this workshop is his work to develop neurocognitive modeling methods. Neurocognitive models fuse together the signal processing methods and machine learning with cognitive models that express probability models of behavior. Most recently, these models have incorporated novel interpretable neural networks to directly estimate cognitive processes from EEG signals.

Dr Srinivasan's talk introduced a framework for combining cognitive models with neural data and AI to better understand decision-making processes. The approach uses variational autoencoders to jointly model EEG data and cognitive parameters like decision criteria. This allows predicting cognitive states from brain data and vice versa. Recent work uses this approach to predict brain signals from behavioral data by learning the relationship between the two. Another experiment he discussed involved a coin flipping task where people had to judge if a coin was biased toward heads or tails. Results showed humans use variable decision criteria and are influenced by early outcomes rather than optimizing for accuracy. Another task involved predicting income from demographic data, where humans and AI models plateau around 80% accuracy separately, but combining their decisions could lead to 100% accuracy due to diverse strategies and information.

Dr Srinivasan highlighted that AI is most useful when information is rich, but humans are relatively more useful when information is poor. This work offers opportunities for human-AI hybridization by inferring cognitive parameters like confidence and trust from physiological signals. This could help optimize how information is distributed to different people and machine-agents based on their individual strengths and strategies. It may also help systems identify when a group has enough information to make a decision, but individuals are hesitating due to confidence issues. Combining cognitive models, brain data, behavioral data, and

AI in a joint framework shows promise for understanding and enhancing human decision making.

2.1.4 Dr Janet Rafner

Dr Rafner is Junior Center Director at the Center for Hybrid Intelligence and affiliated researcher at the Interacting Minds Center at Aarhus University in Denmark. Her PhD research was on “Exploring Human-AI Interaction in Hybrid Intelligence and Creativity through Crowdsourcing and Games,” and her current research focuses on citizen science, participatory futures, psychometric creativity assessment, computational co-creativity, and human–AI interaction. Dr Rafner has received numerous prestigious awards and fellowships, including the Sonophillia Foundation Research Award, a Fulbright Fellowship, a Salzburg Global Fellowship, and the 2023 CircleU Interdisciplinary Research Award.

Dr Rafner’s talk presented current research efforts in the area of hybrid intelligence and human–AI interaction. Her work takes an interdisciplinary approach involving cognitive science, business, and other fields. One of her current projects involves a collaboration with Autodesk that explores developing a hybrid intelligence solution for generative design through interactive ML techniques. She studied concepts like partnership behavior and the willingness of users to contribute to training AI systems. Another project with Adobe involved developing a better interactive ML platform for photo editing. Her group also created a framework to help companies assess where their systems currently fall and how to move toward more hybrid intelligence approaches.

Furthermore, Dr Rafner’s work aims to enable co-creativity between humans and AI. She provided different frameworks for understanding creativity from the perspectives of psychology, human-computer interactions (HCIs), and computational creativity. One of her projects, called “Crea Visions,” aims to allow the general public to envision futures and address problems in their cities using AI-generated images. Through these studies, her group wants to make AI tools that can augment human creative processes like problem identification and solution generation.

2.1.5 Dr Gillian Smith

Gillian Smith is the Director of Interactive Media and Game Development, and Associate Professor of Computer Science, at Worcester Polytechnic Institute. Her research and pedagogy explore human interaction with generative systems in a variety of contexts, including game design, education, and textile arts. Professor Smith teaches courses in novel interface design, artificial intelligence for interactive media, serious games, and creative coding. She also maintains a creative

practice in computational creativity, including in live coding performance, generative textile arts, and game design. She earned her BS in Computer Science from the University of Virginia, and MS and PhD in Computer Science from University of California, Santa Cruz.

Dr Smith's talk provided a comprehensive overview of several key topics related to designing technologies that incorporate AI. She argued that games and media industries have long been at the forefront of AI development and adoption of new technologies. Game designers and developers are frequently early adopters who help drive advances in areas like AI research and even graphics technology.

When it comes to designing AI systems themselves, there is a host of issues related to trust and how people tend to anthropomorphize technologies. Humans tend to ascribe human qualities like intelligence, emotions, and biases to AI even though it is not sentient. This can lead to misunderstandings about a system's actual capabilities. Dr Smith emphasized how the data used to train AI models is important since it can reflect and even amplify the biases of its creators. Unless properly addressed, these biases will then appear in the systems' outputs and decision-making processes.

Dr Smith then discussed her course on the ethics of creative/generative AI. A key part of her course involved critical design projects where students explored potential social consequences of AI development through intentionally speculative and provocative designs. This approach treats design as a valid form of inquiry in its own right. It helps shed light on complex issues by considering alternative perspectives. Some examples of student projects were discussed in depth. One examined gaps between human expectations of culturally familiar concepts and AI-generated outputs, finding representation and linguistic biases. Another looked at diminished concern over increasing AI usage over time through a fictional account. In general, the projects examined how AI-generated content does not always match surface-level appearances and assumptions.

Dr Smith's talk emphasized that while AI offers opportunities, its development and use also raise important considerations around biases, impacts on underrepresented groups, and shifting social norms. Designers play an important role in addressing these issues through reflective practice and envisioning unintended consequences of present technological trajectories. A key take home message from her talk was that "you can't AI your way out of bad design in the first place."

2.1.6 Dr Manish Saggar

Dr Manish Saggar is an assistant professor in Computational Neuropsychiatry at Stanford University and directs the Brain Dynamics Lab. The overarching goal of

his lab is to develop computational methods that could allow for anchoring psychiatric diagnosis into biological features (e.g., neural circuits, spatiotemporal neurodynamics). His lab focuses on developing data-driven computational methods to generate clinically and behaviorally relevant insights from high-dimensional biological data (e.g., neuroimaging) without necessarily averaging the data at the outset. The lab also actively pursues developing novel technologies for experimental design and data collection for enhancing human cognition (e.g., creativity and collaboration). Dr Saggar received his PhD in Computer Science from the University of Texas at Austin and later received postdoctoral training in Psychiatry from Stanford University School of Medicine.

Dr Saggar uses advanced brain imaging and network analysis techniques to study free-form creativity, problem solving, and their underlying brain dynamics. He discussed a study where participants sketched solutions to design challenges in a magnetic resonance imaging (MRI) scanner while having their brain activity, speech, and sketches recorded. Graphs were generated from brain activity over time to analyze connectivity and modularity changes. Low-dimensional embeddings helped identify task-specific signals and cluster similar time points. Comparing graphs from different problem states showed shifts in attention and default mode networks during moments of “inspiration”. Interestingly, these brain graphs would be colored according to states like inspiration, frustration, or typical stages of problem solving, offering an interpretable visual description of the results. This revealed which networks were engaged during each state. Preliminary findings suggested graphs could serve as biomarkers to detect conditions affecting problem solving. For some, atypical brain modularity was seen, suggesting differences in creative thinking. With more data, these methods may help characterize individual differences and clinical populations, with a goal of advancing our understanding of creative cognition and facilitate inspiration by monitoring predictive brain signals. With hyperscanning, AI could provide feedback based on brain states to enhance the creative process. Overall, this multimodal approach could provide novel insights into creativity not possible with conventional analysis.

2.1.7 Dr Stephen Fiore

Dr Stephen M Fiore is Director, Cognitive Sciences Laboratory, and professor with the University of Central Florida’s (UCF’s) Cognitive Sciences Program in the Department of Philosophy and School of Modeling, Simulation, and Training. He maintains a multidisciplinary research interest that incorporates aspects of the cognitive, social, organizational, and computational sciences in the investigation of learning and performance in individuals and teams. His primary area of research is the interdisciplinary study of complex collaborative cognition and the

understanding of how humans interact socially and with technology. He is Immediate Past President of the International Network for the Science of Team Science, and Past President for the Interdisciplinary Network for Group Research. In 2018, Dr Fiore was nominated to Defense Advanced Research Projects Agency's (DARPA's) Information Sciences and Technology (ISAT) Study Group to help the DOD examine future areas of technological development potentially influencing national security. He has been awarded UCF's prestigious Research Incentive Award four times to acknowledge his significant accomplishments. He is recipient of UCF's Luminary Award (2019), as recognition for his work having a significant impact on the world, and UCF's Reach for the Stars Award (2014), as recognition for bringing international prominence to the university. Dr Fiore has also co-authored over 200 scholarly publications in the area of learning, memory, and problem solving in individuals and groups.

Dr Fiore discussed the nature of episodic memory, its role in human cognition, and opportunities for incorporating episodic-memory capabilities into AI. His work highlights the need to develop human-like AI. For example, social cognitive skills distinguish humans from current AI systems. Humans possess episodic memory involving personally experienced discrete events, which in turn are used to define one's sense of self. Episodic memory captures the who, what, where, why, and how of experiences by allowing conscious recollection of past events. Developmentally for humans, around age 3 episodic memory abilities emerge alongside theory of mind, storytelling ability, and a reduction in childhood amnesia. While most AI focuses on semantic knowledge, developing episodic abilities could help machines interpret scenes, understand social interactions, and predict behaviors through theory of mind. Ubiquitous sensors and personal assistants are generating growing sources of episodic data that could be organized into discrete events to form an AI's own episodic experiences. Advances in computer vision, narrative understanding, and other fields provide a scientific foundation.

Military operations could leverage episodic AI through improved mission planning, briefings, and debriefings. By drawing on relevant prior experiences stored in an episodic format, an autobiographical AI system integrated with human teams could enhance knowledge management, preparation, and decision making. This represents an opportunity to develop more human-like social and self-aware artificial intelligence.

2.1.8 Dr Lena Smirnova

Dr Smirnova is an assistant professor at the Environmental Health and Engineering Department at Bloomberg School of Public Health, and Center of Alternatives to Animal Testing, Johns Hopkins University, where she is leading an educational

program on microphysiological systems (MPS). She has joint appointments at Johns Hopkins School of Engineering and Georgetown University. Her research focuses on development of new approach methodologies for developmental neurotoxicity testing and understanding gene environmental interactions in autism. She promotes the idea of organoid intelligence as a new field to address cognition with an in vitro brain model, the brain MPS. She received her PhD from Charite Free University, Berlin, and postdoctoral training at the Federal Institute for Risk Assessment. She is a co-organizer of a series of conferences on MPS and president of the international MPS society.

Dr Smirnova specializes in brain organoids, which are 3D cell cultures that mimic the structure and organization of the brain, containing different brain cell types like neurons, astrocytes, and oligodendrocytes. Organoids open up new possibilities for studying neurodevelopmental disorders, testing drugs, and modeling diseases. Indeed, researchers have grown organoids mimicking conditions like microcephaly, Zika virus infection, and Alzheimer's. Moreover, organoids can be grown from stem cells, allowing the study of genetic variations and modeling of patient-specific conditions. This technology holds promise for developing personalized treatments.

Organoids are relevant for understanding the brain and its computations. Computer models are useful for simple and computational tasks but cannot replicate the complexity and learning ability of the brain. Brain organoids aim to bridge this gap by allowing the input and output needed for learning. Moreover, novel electrode and imaging technologies allow high-resolution recording of electrical activity directly from organoids over long periods of time, advancing our understanding of neural circuit formation and activity. An interesting proposal is an “organoid-AI hybrid” using organoids as a model to test AI systems, with the goal of creating more human-like artificial general intelligence. Inputs to organoids could be provided by AI systems, with organoid responses then processed by AI in a closed feedback loop. Any research combining organoids and AI raises important ethical questions around consciousness, sentience, and potential for harm that must be carefully considered. Researchers emphasized organoids currently lack consciousness and the goal is to combine technologies to better understand the human brain and cognition.

2.1.9 Dr Michael Beyeler

Michael Beyeler directs the Bionic Vision Lab at University of California, Santa Barbara (UCSB). He received a PhD in Computer Science from UC Irvine as well as a BS in Electrical Engineering and an MS in Biomedical Engineering from ETH Zurich, Switzerland. Prior to joining UCSB, he completed a postdoctoral

fellowship in the labs of Ione Fine (Psychology, Institute for Neuroengineering) and Ariel Rokem (eScience Institute) at the University of Washington, where he started working on computational models of bionic vision. He is Associate Director of the UCSB Center for Virtual Environments and Behavior (ReCVEB) and recipient of the National Institutes of Health (NIH) K99/R00 Pathway to Independence Award as well as the prestigious NIH DP2 New Innovator Award.

Dr Beyeler's talk discussed his research on developing augmented sensory technologies and methods for personalizing them to individuals. His Bionic Vision Lab at UCSB is working to create visual neuroprosthetics that can restore or enhance vision for those who are blind or have low vision. Their research utilizes virtual reality simulations to rapidly test and iterate on prototype neuroprosthetic devices. This allows them to model different patients' visual perceptions and evaluate how proposed stimulation patterns may appear to each user. A key focus is developing a preferential Bayesian optimization approach to personalize devices. Through this method, users provide feedback by choosing which of two simulated visual stimuli they perceive as better. This feedback is used by an ML model to iteratively estimate each individual's unique sensory response properties over multiple trials. With a personalized model of a patient's visual system, the researchers can then determine the optimal stimulation patterns required to produce target visual outputs.

The long-term goal is to create an adaptive "smart bionic eye" system. Such a device would tailor visual augmentations and highlights to a user's specific tasks or environment based on computer vision analysis and their contextual awareness. For example, a Soldier's device may emphasize threats, navigation cues, or other mission-critical information. Beyond restoring basic vision, these personalized augmented technologies aim to maximize benefits for applications like enhancing situational awareness. This work aims to develop adaptive neuroprosthetic solutions informed by neuroscience, user input, and ML. Personalizing augmented sensory systems could help optimize their functionality for individuals and applications over time as needs or conditions change. This research may advance technologies to augment human capabilities in impactful new ways.

2.1.10 Dr Sarah Kitchen

Dr Kitchen received her PhD in mathematics from the University of Utah in 2010. She co-leads the AI and Robotics Practice Area at Michigan Tech Research Institute (MTRI) and is a principal investigator for a portfolio of research projects in multi-agent autonomy. Before her current role, Dr Kitchen held a mathematics postdoctoral position at the University of Michigan. Dr Kitchen's research interests

include problems related to intelligence, decision processes, data representations, and foundations for models of multi-agent systems.

Dr Kitchen's talk applied a dynamical systems and optimal control point of view to graph-structured data, in particular vector-valued signals on graphs. An example of such a model is a predator-prey model in which populations can move between disconnected ecological patches. In this case, the patches are assigned to nodes and the population of each species forms the vector-valued signal at that node. The system must account for the population dynamics at each node, via the predator-prey model as before, and must also incorporate the migration between patches along edges of the graph. The underlying graph structure is directly applied to these dynamics through the Laplacian for the graph, which plays the role of the Laplacian in the heat equation. In order to convert this model to a command and control model, we replace the predator-prey dynamics with a combat model called the Lanchester differential equation and allow the commanders to make decisions about maneuver across the edges of the graph.

Effectively, this makes the Laplacian for the graph a function of the commander's decision, and the corresponding model is a bilinear control model. Bilinear control models on networks also appear in engineering and biology applications, such as neurological dynamics models. Finally, an important area of research in reinforcement learning, which may be considered a family of methods for producing approximately optimal policies for very general classes of optimal control problems, is that of state representation. In the implementation of a reinforcement learning method, the true state of a complex environment is usually presented to the learner as an aggregate encoding of features of the environment. Unmodeled attributes of the environment contribute to the stochasticity of the model underlying the reinforcement learner. We propose that graphs and vector-valued signals on graphs are very flexible structures for modeling fine-scale features of a wide variety of environments for reinforcement learning, which can also be aggregated in a principled way using methods of graph theory, dynamical systems, and optimal control on graphs.

2.1.11 Dr Ranjay Krishna

Ranjay Krishna is an assistant professor at the Paul G. Allen School of Computer Science & Engineering at the University of Washington. His research lies at the intersection of computer vision and human computer interaction. This research has received best paper, outstanding paper, and orals at CVPR, ACL, CSCW, NeurIPS, UIST, and ECCV, and has been reported by Science, Forbes, the Wall Street Journal, and PBS NOVA. His research has been supported by Google, Amazon, Cisco, Toyota Research Institute, National Science Foundation, Office of Naval

Research, and Yahoo. He holds a BS in Electrical and Computer Engineering and in Computer Science from Cornell University, an MSe in Computer Science from Stanford University and a PhD in Computer Science from Stanford University.

Dr Krishna's talk discussed several projects at the intersection of AI and human interaction. One project introduced a method of "training robots without robots" using augmented reality to collect robot demonstrations from people interacting with virtual robots in their own homes and environments, allowing for more diverse training data. Another project used computer vision and social interactions on social media to allow an AI system to ask questions to learn about new concepts from people, expanding its knowledge over time. The prompt engineering project evaluated language models by having people use them to solve problems through an interface called "the prompt." This allows assessment of how well people can solve tasks with AI assistance rather than just model performance on fixed tasks. Research on AI-assisted decision making found that explanations can sometimes increase "over-reliance" on the AI and lead to accepting incorrect decisions. By framing decision making in terms of costs and benefits, they found AI support is most helpful for difficult tasks when explanations are easy to understand. His work emphasizes the need to effectively train robots, expand AI knowledge through social interactions, evaluate AI through human problem solving, and provide guidelines for when AI assistance improves or hinders human decision making.

2.2 Army SME Talks: Mission Command Battle Lab

Two US Army SMEs, Mr Thomas Christensen (Director) and Mr Christopher Kelshaw (Chief Analyst, Futures Branch) from MCBL (Fort Leavenworth, KS) accepted our invitation to give presentations each day to educate the workshop attendees specifically on Army C2 processes and lead focused discussions on potential future applications of hybrid human/machine intelligent technologies to specific C2 challenges in this area. We learned that the Army C2 system consists of four subcomponents: Personnel/Organizations, Communications Networks, Processes, and the Command Posts. At the outset of their presentations, Kelshaw and Christensen highlighted four principal capability objectives for US Army C2 to facilitate and focus our discussion and ideation at the workshop:

- 1) Future US Army should leverage machine-aided C2 processes and procedures.
- 2) Future US Army C2 should require relatively fewer people and organizations.

- 3) Future US Army C2 should allow commanders to make faster, less error-prone, and more insightful decisions than the enemy.
- 4) Future US Army C2 should function well, even when command post nodes are dispersed and, at times, moving.

To summarize their perspective on future US Army C2, there is an emerging and lethal danger to command posts as they have been implemented in past and current-day US Army operations. With the increase of technology, especially with a near-peer adversary with advanced sensing and intelligence gathering capabilities, large command posts (in terms of physical size and personnel) are too easily detected and targeted. We must, therefore, adapt and reenvision future command post structures to increase survivability. The prevailing idea is that future command posts should be 1) smaller in physical size, 2) smaller in terms of the number of personnel, 3) agile with quick repositioning capabilities, and 4) dispersed into many smaller nodes with robust intra-, inter-, and extra-nodal communication network capabilities.

With this in mind, workshop attendees were challenged to brainstorm and consider how future technologies could aid in meeting these objectives. C2 is critical for integrating and synchronizing all warfighting functions toward a common operational purpose such that winning a battle is improbable without effective C2. Despite the importance of the C2 system and the clear imperative that it must change in the future to be both survivable and functional, it is still unclear exactly how to get there. This creates a major challenge and opportunity for researchers and technologists to develop new lines of research and tools to radically change and enable C2 capabilities in the future Army.

We believe ARL is well-positioned to take a leading role in research and development of these future technologies and capabilities. While the objectives outlined previously are clear, the potential solution space is vast and will require collaboration and coordination among many groups including academic researchers with a diverse range of expertise, Army SMEs, concept developers, and more. The Hybrid Thinking Workshop can be seen as a first step forward in this direction, seeding critical future US Army C2 objectives in the minds of ARL and academic researchers and bringing together a diverse group of scientists to begin ideating and collaborating in this critical research area alongside Army SMEs.

2.2.1 Day 1

The presentation by MCBL on Day 1 of the workshop provided a broad overview of the C2 system, focusing on people and processes. Because war is dynamic in nature, it is important for commanders and their staff to continuously update and

refine their mental models of the operating environment as operational and mission variables change. This process must involve communication from the commander to their staff to create a shared visualization and communication from the staff back to the commander providing information to either support or change the current visualization. Development of a shared vision is critical for this process to run smoothly and effectively because, in many ways, the staff serves as essentially “an extension of the commander’s brain.” This is relevant to modern and future warfare because the scale of the complexity is growing beyond the comprehension capabilities of any one person, or even an entire professional staff.

The commander’s staff in the present-day C2 system can be thought of as specialized *knowledge workers*, with deep but narrow expertise and responsibilities aligned with their role in the system at large. Typical duties of a staff officer include advising and recommending, collecting information, developing running estimates, assessing operations, and preparing plans/orders. Staff are organized into functional cells including intelligence, movement and maneuver, fires, protection, sustainment, and other staff sections, as well as integrating cells responsible for integrating information across the various functional cells for the purposes of managing current operations (CUOPS), future operations (FUOPS), and plans. The time horizon is typically days-hours for CUOPS, weeks-days for FUOPS, and months-weeks for plans, but all three of these integrating functions happen simultaneously and continuously on the battlefield such that plans transition products to FUOPS, who transition products to CUOPS as the mission unfolds.

Due to the increase of technology and specialization areas, as well as organizational inertia, the command staff has grown substantially over time. At the division level, for example, there are approximately 500 staff personnel involved in the C2 system. In many cases, there are staff in different functional cells that have similar duties and responsibilities but with a different emphasis. As such, there was discussion during the workshop about the potential impact of developing future intelligent machine-aided technologies to support a fundamental activity that is shared across many different staff roles, such as automating routine workflows or writing support. Future research could be helpful in identifying these common processes as a target for technological innovation as a means for making C2 processes more effective and efficient, and potentially meeting the objective of reducing staff size.

The increasing complexity of the battlefield is another challenge faced by commanders and is a ripe target for future intelligent machine-aided technologies. Operational variables that must be considered during Army operations include political, military, economic, social, information, infrastructure, physical environment, and time. Moreover, mission variables must also be considered including the mission, the enemy, terrain and weather, troops and support available,

time available, and other civil considerations. All variables must factor into the decision-making process, and intelligent technologies could play an outsized role here due to the speed, scale, and computational power of machines. Complexity could be ameliorated with technology that enhances complex visualizations, organizes information, assesses relevancy and complex dependencies, and performs complex reasoning relative to objectives and the desired end state. Time is a critical factor in this process because there are often only brief windows of opportunity to exploit an enemy's vulnerability or provide protection by disrupting enemy actions. Intelligent machines could, in principle, help commanders to recognize emergent windows of opportunity and risks more robustly and more quickly from the myriad of complex and dynamic information from the battlefield.

The C2 system comprises three processes: operations, integrating, and planning. The operations process itself is constantly ongoing and involves a continuous loop of planning, preparing, executing, and assessing. The first three stages of the operations process require continuous assessment to ensure that the operation is unfolding optimally to meet the objective or desired end state; otherwise, adjustments must be made in the following cycle of plan -> prepare -> execute. Integrating processes are found within the operations processes. These include intelligence preparation of the battlefield (IPB), information collection (IC), targeting, risk management, and knowledge management.

Some integrating processes are also ripe targets that could be dramatically improved by intelligent machine-aided technologies. As an example, we have vast new information collection technologies from satellites, drones, and novel sensors, and so on, and machines could play a critical role in analyzing, integrating, and interpreting those raw signals for battlefield-ready insights to inform a commander's decision making. Another example is related to knowledge management—the integrating process that connects everything by characterizing how knowledge is organized and shared across the staff. Information flow can be characterized, in part, through the battle rhythm, a carefully managed linear sequence of meetings and working groups composed of different elements of the staff. Each unit and mission type requires a bespoke battle rhythm to effectively share knowledge. Intelligent machine-aided technologies could be designed to adaptively optimize battle rhythms for a given staff structure and mission or could be used in novel ways to reduce the reliance on battle rhythms altogether to share knowledge continuously and instantly across units and across echelons. Technologies that support rapid, dynamic visualization and sharing of information could obviate the need for sequential, discrete meetings to exchange information.

Another serious consideration related to knowledge management is the rapid turnover rate of staff roles such that commanders and staff officers often spend just

1 to 2 years in each role before moving or advancing to a new position as their career progresses. A more immediate challenge associated with knowledge management during an operation is the daily shift change from one staff officer to the next for rest, injury, and so forth. Intelligent machine-aided tools that can aid in the transfer of both tacit and explicit information and knowledge about the present battlefield situation between individuals more effectively and efficiently could be extremely helpful in ensuring there is not a lapse of operational tempo or loss of operational information following shift changes.

Plans are often more conceptual and far removed from events on the battlefield, so the military decision-making process (MDMP) is used to translate plans into coordinated action. MDMP includes the following processes, in order: mission analysis, course of action (COA) development, COA analysis (wargaming), COA comparison, COA selection/approval, and orders production. Machines could conceivably aid MDMP in multiple areas, especially by providing operationally relevant modeling and forecasting capabilities to visualize various ways in which the operation could play out. In particular, COA development is a relatively long process (though in practice it is often cut short due to time pressures) that could benefit from intelligent machine-aided technologies that understand Joint and Army doctrine, operational goals, mission variables (i.e., terrain and weather considerations), intelligence (enemy force positions and red team capabilities), blue force capabilities, logistics, and so on, to generate realistic, or even creative, COAs for the staff to analyze and consider. Further leveraging explainable AI techniques could provide additional context for recommended COAs from intelligent COA development tools.

Decision support tools are developed during the MDMP process (during COA development and analysis, that is, wargaming). The decision support matrix is a document (a table) that lists anticipated critical events or “decision points,” Commander’s Critical Information Requirements (CCIRs) that are tied to these decisions, an event description, friendly action, and location information. During the mission, staff functional and integrating cells are responsible for tracking events during the operation and determining whether a decision point’s associated CCIR criteria have been met and if an adjustment is needed as specified through if-then conditional statements. The decision support template has similar functionality, but instead of information contained in a table it is pictorially represented on a map of the area(s) of operation (AO). Intelligent machine-aided tools that automatically provide timely updates of relevant battlefield information to ascertain specific decision points could be a very useful and impactful tool to ensure that critical information is not missed and is provided as soon as available.

The common operating picture (COP) is a visual representation that provides a visual depiction of the AO for monitoring and managing mission execution and ongoing planning. It provides relevant information to C2 personnel/organizations to determine whether mission execution is successfully matching their currently held mental model and objectives and allows staff to share in a common visualization to improve shared situational awareness. Joint, integrated, and effective fighting forces require a shared visualization and, therefore, an effective COP that includes aspects of and effects in all domains—land, air, maritime, space, and cyber. Understandably, maintaining a shared and dynamic COP will be an increasing challenge for future command post nodes that are more dispersed, as the same information may not be available to all command post nodes at the same time. This is a fundamental challenge that must be solved to successfully integrate and synchronize Joint and Army warfighting functions and multi-domain effects in the future; failure to plan and act with a shared COP could be the principal cause of real-world mission failures. Moreover, each commander may have a different preference for what, where, how, and when to visualize certain information in the COP representation, and all staff need to be on the same page. Future intelligent machine-aided tools could help to populate, synchronize, standardize, and adapt information that is presented in the unit’s COP at any point in time during the battle to streamline and optimize decision making according to the needs or preferences of the commander.

In summary, day 1 of the workshop provided a detailed overview of the C2 system, including the people and processes involved. It provided context and specific areas of interest for considering where research efforts could be invested to impact C2. It laid out several objectives to work toward, including how to empower commanders to make better, faster decisions with less people that are distributed and may be on the move.

2.2.2 Day 2

The focus of the MCBL presentation on Day 2 of the workshop was a deeper dive into specific C2 processes, with a focus on integrating processes. There is an expectation that the future of warfare will compress the time frames associated with CUOPS, FUOPS, and planning, so there is a need and opportunity to introduce new technologies to make these processes more agile and timely. Effective planning is a critical factor in synchronizing and integrating our forces to shape and control the battlefield for mass effects across multiple domains (cyber, space, land, air, sea). An interesting idea was discussed that “all plans are useless, but planning is essential,” which captures two important notions. The first is that warfare is complex and unpredictable (the enemy always gets a vote), so expecting events to

unfold according to a detailed plan is unrealistic and futile. Plans must continuously be adapted and adaptable, which conflicts with the core notion of a *plan* as being a sequence of action to be implemented. Secondly, the in-depth process of *doing* planning has value in-and-of itself that is beyond the end product of producing a plan. The planning process provides an invaluable tool for staff to play out different scenarios and gain a better situational understanding of the challenges and opportunities that could arise during an operation. If highly automated, intelligent machine-aided planning tools are developed and adopted as part of future C2 process execution, there is a potential risk of staff over relying on the automated tools and failing to derive important learning from processes themselves.

IPB is an aspect of MDMP that characterizes the threat in the environment including who, what, where, when, and how the threat may behave during an operation. Most of the staff is involved in IPB, and it is a sequential task that results in information overlays (whether acetate film with markings or, more recently, Microsoft PowerPoint slides) on maps and an event matrix (a table) that represents the shared vision of an operation including named areas of interest and threat decision points. IPB analysis involves analyzing terrain (avenues of approach, mobility corridors), the physical locations, composition, and strength of the enemy in operations, enemy doctrine and preferred tactics, the distance/time between threat forces and control measures, key weapons systems, high-value targets, and more. This information is represented as an event template that serves as a guide for collection planning that identifies threat decision points, indicators of threat activity, time phase lines, and named areas of interest. IPB was identified as a prime target for applications of machine intelligence and automated tools including those that (i) help understand complex terrain and model effects on friendly and threat operations, 2) allow accurate inference of threat positioning and intentions from incomplete information, and 3) enable realistic modeling and visualization of threat plans over time.

IC is a process that defines how we collect relevant threat information based upon IPB outputs to help inform commander decision making. There are numerous tactics and technologies to support information collection from satellite imagery and unmanned aerial vehicles (UAVs) to Soldiers conducting reconnaissance on the ground. This process results in an information collection matrix (a table) that contains information about priority intelligence requirements, indicators, specific information requirements, start and stop times for collection, and assignment of friendly assets to perform the specified collection activities. Intelligent machine-aided tools that process information from multiple sources or sensors and improve data visualization of battle dynamics could be an extremely helpful planning tool for information collection.

The targeting process is guided by the information collection outputs and determines when, where, and how we plan to target and destroy enemy units, including developing a list of prioritized targets (command post nodes, sustainment units, high payoff targets, etc.). It involves four stages including decide, detect, deliver, and assess (D3A). Potential technology focus areas to impact the targeting process include a system to propose targeting priorities based upon a system-of-system approach and tools that enable the creation of sensor coverage, sequencing, targeting options, and reprioritization of targets when necessary.

Knowledge management was described as the backbone running through all other processes by ensuring that the right information is given to the right person at the right time in the right form and in the right place. Battle rhythms are a centerpiece of knowledge management and serve to synchronize and sequence information sharing events (meetings, working groups, etc.) among the staff cells and the commander. The battle rhythm is carefully sequenced such that each event has specified inputs and outputs that serve as inputs for subsequent events, all to help commanders make timely and informed decisions over time. We discussed that the current approach to knowledge management is reflective of the industrial era (i.e., linear, centralized, hierarchical), and the need to develop new tools and processes to transition to the knowledge era (i.e., decentralized, concurrent, agile). The current approach to knowledge management may be effective but is inefficient in that it relies heavily on face-to-face (or virtual) meetings of two or more people discussing information with a purposefully limited scope and requires staff to prepare specific subsets of information (often in the form of Microsoft PowerPoint slides) targeted for the audience and meeting topic. In this way, intelligent machine-aided technologies could play an outsized role in improving knowledge management in ways that impact all other C2 processes. One key idea is the need for information to be shared/updated continuously and more broadly via live dashboards that are accessible across staff elements and across echelons. The nature of knowledge management would thus become less sequential and linear, and information could be accessed universally, as needed, much like the way a smart phone in your pocket gives you access to near unlimited information accessible via the internet at the point of need.

COA development, as part of planning, is a rather time-consuming, but a fundamental, planning process step that includes analyzing relative combat power, generating options, arraying initial forces, developing concepts of operation, assigning responsibilities, and preparing COA statements and sketches (visualizations). Future technology focus areas to assist COA development include systems that enable calculation of multi-domain relative combat power and correlation of forces and means, and intelligent machine-aided generation of

statements and sketches. There is an imperative to both speed up the COA development process because time constraints often prevent consideration of multiple COAs during combat operations and enable more creative COAs to be developed and considered. Developing more than one COA reduces risk by ensuring that multiple options can be created with the best of these ultimately selected. More diverse creativity in COA development could unlock unconventional (and therefore unpredictable) operational approaches that outperform default or common COAs developed quickly under time pressure.

The COA development process can be thought of as creating a set of hypotheses to solve a military problem that are then tested in the COA analysis, or wargaming, process. Wargaming involves a sequential, cyclical process of simulating friendly and threat actions, reactions, and counteractions through the specified COA. The wargaming process itself is a critical tool for collective staff learning, developing a shared visualization, and ensuring that the developed COA is rigorously tested. This is another area identified in which intelligent machine-aided technologies could deliver massive improvements over current methodologies; for example, systems that enable rapid simulation and assessment of multiple (perhaps thousands) wargaming iterations over multiple performance criteria to optimize COA analysis. In the ideal case, IPB would generate three potential enemy COAs, and staff would create three potential friendly COAs. All COAs would then be wargamed against each other. However, during actual operation execution there is often not enough time to do this completely solely by humans, so wargaming simulation tools could be extremely helpful to fill the gap, especially under time constraints.

In summary, day 2 of the workshop included a more detailed presentation and discussion of various integrating processes and MDMP, leading to fruitful discussions about specific technology focus areas and intelligent machine-aided technologies to improve activities in the C2 system.

2.2.3 Day 3

The focus of the presentation by MCBL on day 3 of the workshop involved characterizing the patterns of work for staff officers to better understand common tasking and cognitive demands on individual staff members performing their duties as cogs in the larger C2 system. Common staff member duties include the following:

- Advising and informing their commander
- Building/maintaining running estimates
- Providing recommendations

- Preparing plans, orders, and other staff writing
- Assessing operations
- Managing information in their area of expertise
- Identifying and analyzing problems that arise
- Conducting staff assistance visiting
- Performing risk management
- Performing IPB
- Conducting research
- Performing administrative procedures
- Exercising staff supervision

As can be seen in the list, staff members are very much like “knowledge workers” in that their duties require research and analysis to distill insightful and/or actionable information from disparate sources. Communication skills are highly prized in an effective staff officer, which involves managing not just information in their area of expertise, but also managing and interacting with people. This social aspect can introduce additional cognitive and emotional stressors to impact performance of their duties. Other valuable skills include being able to visualize and analyze data and communicate those insights to other people verbally, visually, and in written formats.

Army C2 consists of both command elements and control elements. The latter includes advice and feedback provided by staff members to the commander in the form of running estimates. These are maintained by staff functional and integrating cells in various forms across the warfighting functions (IPB example: How many air defense systems does the threat currently have in this area? Logistics example: How much fuel and ammunition is remaining or will be required for an operation?) to aid planning and decision making. Running estimates could greatly benefit from automated sensors feeding a system with a live dashboard to help maintain an accurate operational picture and reduce the guesswork involved in estimating various factors (reserve fuel, ammunition spent, etc.). Intelligent machine-aided tools for staff writing and reports (that often follow a template structure) could be extremely useful to free up time to focus on other more creative aspects of C2 process activities. Technologies to automate routine workflows associated with multiple staff roles could help to reduce organizational size and allow staff to focus on higher-order reasoning and more cognitively demanding tasks.

Another technology focus area discussed is the development of an AI copilot or personal assistant with episodic memory capabilities that continuously learns from experiential data, much like humans do. An AI personal assistant could be tailored to a user and their position and provide insightful, specific feedback to the user as almost an extension of their brain to be queried as needed.

A discussion followed about the way in which information is typically shared among C2 personnel and organizations, identifying that human–human communication is the predominant method (i.e., face-to-face conversation or radio conversation) rather than interacting with machines (i.e., accessing information directly from a computer interface). In a future command post structure that is dispersed into smaller nodes that may be moving and out of communication for periods of time, new technologies should be developed to maintain adequate information sharing without requiring human–human contact such as access to universally accessible, secure, and up-to-date information systems so relevant data can be pulled and analyzed by remote staff members as needed.

Finally, we discussed the use of future collaboration and communication technologies to support telepresence and enable commanders and staff officers to interact remotely while maintaining social presence and situational awareness. A technology focus area should be developing novel forms of social interaction, perhaps via augmented reality or virtual reality (AR/VR) headsets or other modalities, to foster communication and information sharing while individuals are disconnected and/or on the move to evade threat targeting.

3. Conclusion

The battlefield will continue to evolve at a rapid pace with the introduction of complex, intelligent technologies by The US Army and its adversaries. We envision that future hybrid technologies will help ensure that US Army Soldiers are prepared for this complexity by leveraging the synergistic strength of Soldier intelligence and machine intelligence. We anticipate that such technologies will enable faster adaptation to technological and environmental changes, more robust planning and real-time understanding of battlefield events for mission command, and faster decision making to exploit ever-shrinking windows of opportunity to gain a decisive edge on the battlefield.

To this end, we organized the Hybrid Thinking Workshop to help kick start a new research program focused on developing hybrid intelligent human–technology capabilities to support the future Army C2 system. With this goal in mind, we brought together an eclectic mix of academic researchers with Army C2 SMEs from MCBL to foster dialogue, information sharing, and collaborative ideation for three

days at Stanford University. The design of the workshop enabled information sharing through a series of research presentations from invited academics, and information presentations from ARL researchers and MCBL concept developers. The breakout sessions in the afternoon allowed smaller groups (8-10 people) to engage in focused brainstorming to develop creative ideas for future concepts of intelligent machine-aided systems to address Army C2 challenges. Below we aim to summarize some of the key points and core ideas from those discussions related to the original capability objectives outlined by MCBL at the start of the workshop.

3.1 Future US Army C2 Should Require Relatively Fewer People and Organizations

As stated, the current C2 system is bloated with personnel; the Division level, for example, has roughly 500 Soldiers in the command staff. As envisioned, future combat operations will require more distributed and agile command post nodes and less forward-exposed units to increase survivability and functionality of the C2 system. How can future intelligent machine-aided technologies support this goal?

We discussed the need to analyze commonalities of staff duties across echelons and warfighting functions and identify repetitive tasks and workflows that could be automated or supported by hybrid human-machine technology. As an example, a core job of many staff officers, regardless of the expertise area, is to analyze, synthesize, and select information to be shared with the commander. Intelligent systems could aid in collecting, analyzing, sorting, prioritizing, and visualizing data for staff members to help them quickly access and summarize relevant information for specific working groups or meetings.

Writing is another common duty of staff officers (orders production, etc.), and much writing follows a prespecified structure or template. Automated tools to help generate written material could free up valuable time for staff to focus on more challenging cognitive tasks and higher-level reasoning. Together, we believe that targeting research and technology capability development to support or automate routine tasks and workflows encountered by most staff will free up staff time and resources to either get more done with the same staff size or achieve comparable performance with a smaller staff size.

Another possible avenue toward reducing staff size is to completely envision the MDMP process itself. A take-home message from one of the academic attendees, Dr Gillian Smith, was that “you can’t AI your way out of bad design,” and much discussion during the breakout sessions centered on “Why does the Army do it this way?” and “Why doesn’t the Army do it that way?” It is beyond the scope of this workshop and this report to advise the Army on exactly which large-scale changes

to the C2 system could reduce staff size and make planning and decision making faster. However, with the continuous development of new intelligent technologies, there is an opportunity for the US Army to revamp the military decision-making and planning processes from the ground up. Right now, the process is linear, analytical, and time-consuming—and the Army has been training and executing C2 in roughly the same way for over 5 decades. Intelligent machines and advanced computing systems, in collaboration with human teams, could help flatten the C2 organization and decision-making process to make it more parallel, fast-paced, and responsive to complex and dynamic events.

3.2 Future US Army C2 Should Allow Commanders to Make Faster, Less Error-Prone, and More Insightful Decisions Than the Enemy

One of the biggest challenges for commanders is they do not have time to gather all (or even most) of the relevant information in time for decisions. We believe that Army commanders could benefit from intelligent machine-aided decision support tools to enable faster and less error-prone decision making. Future intelligent systems designed to continuously analyze real-time battlefield information for context, insights, risks, and opportunities associated with friendly and/or threat actions would help a commander and their staff maintain peak situational awareness to make informed decisions on up-to-date information.

A core contribution of future intelligent tools could be information triage, by collating and filtering voluminous data and organizing and formatting it to help staff focus on what information is most important and relevant at a particular moment in time. We discussed the development of accessible databases and dashboards to host myriad information in a format that is context-aware and simple to query for decision-making support at the moment of need. The notion is to think of information sharing across the C2 system less as “pushing” (as is currently done when staff push information to the commander), and more toward “pulling” information dynamically, as needed, from computer systems to support decision making.

Decision making in MDMP is also slowed by human cognitive and physical constraints that could be overcome by machine-aided technologies. COA analysis, or wargaming, is a time-intensive process that requires a group of staff members to iteratively simulate action-reaction-counteraction cycles of potential battlefield events—a process that is slow and cognitively burdensome. Wargaming could be supplemented, or even radically changed, by intelligent modeling and simulation tools that can take into consideration multiple operational and mission variables

(terrain, enemy positions and capabilities, weather, threat doctrine, etc.) and generate distributions on the likelihood of potential events or outcomes given a particular course of action. Such simulation tools could speed up the COA analysis process substantially while providing novel insights on the likelihood of a sequence of events following a particular action or decision to optimize COA selection.

3.3 Future US Army C2 Should Function Well, Even When Command Post Nodes Are Dispersed and, at Times, Moving

Hybrid intelligent technologies could support future C2 system concepts by enabling continuous and adaptive coordination of distributed units across multiple echelons under conditions of limited or no communications. Current concepts of knowledge management rely heavily upon person-to-person communication—the battle rhythm, for example, codifies the commander’s schedule that entails a linear sequence of meetings with various specialized working groups to gradually develop a complete visualization of the battle and operational environment over time. A future with dispersed command post nodes and limited communication will not easily support the present model of knowledge management and decision making in the C2 system. Instead, we should aim to develop technologies to support information sharing without requiring interpersonal meetings of working groups that provide narrow updates as inputs to yet more working group meetings.

We envision the future command staff roles as shifting toward primarily that of a technology “operator,” with broad expertise in operating and controlling intelligent systems. Such systems should enable universal, secure access to up-to-date battlefield data that has been interpreted and distilled by human-intelligent-agent teams to generate actionable knowledge and insights about the battlefield in the context of the current operation. This knowledge should be available to the commander at the push of a button and intelligently visualized in information displays (map overlays, sorted tables and charts, etc.) so the commander can dynamically access all relevant information during planning and replanning without waiting to attend meetings of working groups—meetings that may or may not even be feasible in a future C2 system with sparse and highly distributed nodes.

Future intelligent technologies should also play a role in ensuring that the “common operating picture” is indeed common and shared across dispersed command post nodes. In present-day Army operations, if a command post loses contact with other command posts and information is not adequately shared, then there is a risk of each node having a dramatically different operating picture due to incomplete, noisy, or out-of-date information. A central COP should be maintained continuously by human-intelligent-agent teams and distributed to all lower echelon

commanders and command post nodes as soon as communication is available and secured. We should develop technologies that enable a universal COP that is comprehensive, user friendly, and constantly maintained and shared across the commanders and their staffs.

3.4 Priority Research Questions Going Forward

The Hybrid Thinking Workshop offered a unique opportunity for ARL researchers to engage with top academics from a variety of research areas and Army C2 concept developers to discuss emerging advances in machine intelligence and human neuroscience and ideate on potential future applications of hybrid technologies to enhance and support military decision making in the C2 system. ARL's Hybrid Human-Technology Intelligence research program will benefit from these discussions and interactions as the research portfolio develops and matures in this exciting new research area.

Several priority research questions were identified that emerged from discussions at the workshop.

Basic research questions:

- What is the definition of hybrid intelligence by comparison to earlier waves of AI? What are its key distinguishing characteristics?
- How do we validate, understand, and check AI or hybrid human–AI systems effectively, and can we identify a taxonomy to describe when/how we should do so in complex socio-technical systems?
- Can we do rigorous science to introspect on foundational models to better understand when they will or will not work in a given application or domain?
- Humans have compressed and adaptive mental models of the world (i.e., intuition) by comparison to machines, but machines have a high capacity for recursion to branch further in deterministic settings (i.e., chess). Where are the boundaries and conditions for determining which approach would be better and associated metrics?
- What is the critical threshold that triggers a person to go from confident to not confident in their decision making, and how does this change for group decision making during different phases of team dynamics?
- How can we effectively probe and capture intrinsic assumptions and operational variables individuals are using in their decision processes for

complex problems, and how do we apply what we learn to human-machine teaming?

- How will our understanding of expertise change as integration of AI tools enable the future average-level human performance to approximate today's peak level of expert performance, and what are the risks and benefits associated with this in different domains?
- Can we achieve expert-level performance without expert-level skill and understanding? Should we? How do we define "understanding" and "expertise" when everyone is performing at the same capacity?
- What risks are there to a future "domain expert" that has gained expertise in using a hybrid technological tool but has failed to gain foundational expertise about the problem space? How much worse will this expert perform if the technological tool fails and they are left to their own devices?
- Everything an individual person knows starts as an episodic memory, which is very different from how computers learn. What are the team/group/organizational concepts of episodic memory, and how do we define and/or leverage episodic memory for human-AI teams?
- How can we use AI to prompt synthesis of new ideas as a truly creative partner? What is needed to enable human-centric hybrid co-creativity?
- How do we ensure that hybrid intelligence accepts the intrinsic messiness of human experience, behavior, and thought (i.e., cognitive biases, decision processes, history and learned behavior patterns, individual concepts of self and personal identity, physical capacities for memory and learning)?

Applied C2-related research questions:

- How can we flatten the C2 organization and decision-making process to make it more parallel, fast-paced, and responsive to complex and dynamic events?
- How can we identify and automate repetitive tasks and workflows that are common across command staff roles?
- How can hybrid technologies improve COA generation and analysis by making the process faster, more exhaustive, and more creative?
- How can automated or hybrid AI tools analyze raw or processed data from multiple domains to identify critical information/events to support adaptive decision making to rapidly exploit windows of opportunity?

- What kinds of mission planning capabilities have been deemed impossible but may be newly enabled by hybrid intelligence?
- Thrust areas are to increase adaptability, creativity, and speed—how are these concepts interrelated and how are they distinct in the context of Army C2?
- Can we design hybrid intelligence technologies that ensure that the right information gets to the right person at the right time and in the right format to streamline and optimize information flow in the C2 system?

List of Symbols, Abbreviations, and Acronyms

AI	artificial intelligence
AO	area(s) of operation
ARL	Army Research Laboratory
C2	command and control
CCIR	Commander's Critical Information Requirements
COA	course of action
COP	common operating picture
CUOPS	current operations
D3A	decide, detect, deliver, and assess
DEVCOM	US Army Combat Capabilities Development Command
DOD	Department of Defense
EEG	electroencephalography
fMRI	functional magnetic resonance imaging
FUOPS	future operations
FY	fiscal year
HCI	Human-Computer Interaction
HHTI	Hybrid Human-Technology Intelligence
IC	information collection
IPB	intelligence preparation of the battlefield
ISAT	Information Sciences and Technology
MCBL	Mission Command Battle Lab
MDMP	military decision-making process
MEG	magnetoencephalography
ML	machine learning
MPS	microphysiological systems
MRI	magnetic resonance imaging

MTRI	Michigan Tech Research Institute
NIH	National Institutes of Health
ReCVEB	UCSB Center for Virtual Environments and Behavior
SME	subject matter expert
UC	University of California
UCF	University of Central Florida
UCSB	University of California, Santa Barbara

1 DEFENSE TECHNICAL
(PDF) INFORMATION CTR
DTIC OCA

1 DEVCOM ARL
(PDF) FCDD RLB CI
TECH LIB

2 DEVCOM ARL
(PDF) FCDD RLA FB
S THURMAN
J GARCIA