Emerging Innovations for Next Generation Mission Planning and Debrief

Kevin A. Gluck, Joshua Ziegler, Brandon Nolan, AaSHAE Eberle, Donald Duckro, James Cline, Luke Nelson, Veda Setlur, Zachary Wallace

Air Force Research Laboratory, L3 Technologies, Rickard Consulting Group

Wright-Patterson AFB, OH

kevin.gluck@us.af.mil, joshua.ziegler@us.af.mil, brandon.nolan.2@us.af.mil, aashae.eberle@us.af.mil, donald.duckro.2@us.af.mil, james.cline.14.ctr@us.af.mil, luke.nelson.12.ctr@us.af.mil, veda.setlur.ctr@us.af.mil, zachary.wallace.9.ctr@us.af.mil

ABSTRACT

The pointy ends of our national defense spears involve some of the most sophisticated technologies in the world. Just behind these pointy ends, however, is a collection of planning and debriefing processes mired in decades-old technologies (e.g., chalkboards, whiteboards, laminated maps) that impede the operational ability to adapt, decide, and execute. This paper describes our research, development, and evaluation efforts for next generation mission planning and debrief. The objective is to simultaneously accelerate processes and improve outcomes. Our approach involves combinations of new technologies that allow human-machine teams to engage in these iterative processes together. The core of our in-house development is a web-based system we call Metis, after the ancient Greek goddess of wisdom, prudence, and deep thought. Metis provides a web interface for people, an API for agents, and a database backend to serve both in real-time. Using this service-oriented architecture, we have begun to develop and integrate an array of agent technologies. These include an air tasking order parser, a reading agent that extracts constraints from mission planning guidance documents, a mission plan generation and validation agent, a planning product development agent, and an automated debrief focus point identification agent. Each of these and their integration is in the prototype stage of development, but already we have begun formative evaluations. These indicate dramatic improvements in workflow and resulting decreases in process completion time. Early comparisons of machine-generated content to human-generated content show that they are often (but certainly not always) comparable. The paper describes all of these components and the latest evaluation results in detail.

ABOUT THE AUTHORS

Kevin Gluck, PhD is a Principal Cognitive Scientist at AFRL, where he serves as the Core Research Area Lead for Personalized Learning and Readiness Sciences and as the Mission Planning and Debrief Technical Lead.
Joshua Ziegler is an Associate Computer Scientist at AFRL, where he supports efforts to improve productivity in cognitive model and agent development.
Brandon Nolan is an Associate Computer Scientist at AFRL, where he serves as a software developer/co-lead developer for Mission Planning and Debrief and as a software developer for the Synthetic Teammate Project.
AaSHAE Eberle is a program manager at AFRL, performing program management functions for the Autonomy Capability Team 3 (ACT3) and formerly for the Mission Planning and Debrief program.
Donald Duckro, PhD is a Senior Programs Engineer at AFRL, where he serves as Program Manager for Warfighter Readiness Research.
James Cline is a Software Engineer at L3 Technologies, where he serves as a Mission Planning and Debrief Software Lead.
Luke Nelson, PhD is a Software Engineer at L3 Technologies, where he serves as a lead programmer and robustness evaluator for Mission Planning and Debrief.
Veda Setlur is a Principal Software Engineer at L3 Technologies. Her areas of expertise are web development, cyber security, and systems engineering.
Zachary "Zen" Wallace is a subject matter expert on 5th generation fighter tactics and current operations with Rickard Consulting Group, where he contributes to the Air Force's training research and technology development.

Emerging Innovations for Next Generation Mission Planning and Debrief

Kevin A. Gluck, Joshua Ziegler, Brandon Nolan, AaSHAE Eberle, Donald Duckro, James Cline, Luke Nelson, Veda Setlur, Zachary Wallace

Air Force Research Laboratory, L3 Technologies, Rickard Consulting Group

Wright-Patterson AFB, OH

kevin.gluck@us.af.mil, joshua.ziegler@us.af.mil, brandon.nolan.2@us.af.mil, aashae.eberle@us.af.mil, donald.duckro.2@us.af.mil, james.cline.14.ctr@us.af.mil, luke.nelson.12.ctr@us.af.mil, veda.setlur.ctr@us.af.mil, zachary.wallace.9.ctr@us.af.mil

INTRODUCTION

The pointy ends of the United States' national defense spears involve some of the most sophisticated technologies in the world. Just behind these pointy ends, however, is a collection of planning and debriefing processes mired in decades-old technologies (e.g., chalkboards, whiteboards, laminated maps). These media and methods were state of the art prior to the information age. They remain sufficient as long as the threat level is relatively low, the mission relatively simple, and the time pressure relatively weak. However, that is not the future the DoD is anticipating. The National Defense Strategy (2018) and the USAF Science and Technology Strategy (2019) make clear the requirement that we be ready to face capable foes in complex and time-compressed operations. They call for accelerated innovation and technology transition to improve our operational ability to adapt, decide, and execute faster and more effectively than ever before. We absolutely cannot be satisfied with business as usual.

Most of the investments in game-changing new capability focus on effects delivery during mission execution. This is understandable, given the criticality of those moments. However, this focus ignores the broader mission cycle within which mission execution takes place (Figure 1).

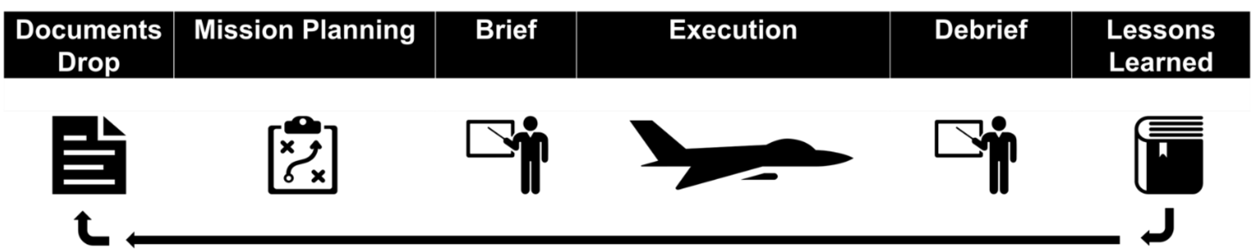

Figure 1. The mission cycle.

The structure and core processes associated with this cycle have emerged as best practices for winning military engagements. Thorough mission planning and briefing are necessary to ensure alignment with intent and a shared representation of how to deal with contingencies. Thorough debrief is necessary to ensure as much as possible is learned from the completed mission, both for the benefit of the warfighters involved and also in order to feed those lessons into future mission preparations.

As mission complexity and consequence increase, so too do the combinatorics and timelines associated with preparing for and learning from mission execution. There is a great deal of competitive advantage to be achieved through improvements to the effectiveness and efficiency of planning and debriefing in future operations. The objective is to simultaneously accelerate processes and improve outcomes. Here we describe our emerging innovations in this area.

TECHNICAL APPROACH

Our general approach is a service-oriented architecture that scales for net-centric interaction with a variety of software agents and people, agnostic to whether they are local or geographically distributed. This involves combinations of new technologies that allow human-machine teams to engage in these iterative processes together (Figure 2).

Metis

The core of our in-house development is a web-based system we call Metis, after the ancient Greek goddess of wisdom, prudence, and deep thought. It supports distributed, real-time collaboration among people and artificial agents. Metis provides a user interface (UI) for people, an application programming interface (API) for agents, and a database backend to support both in real-time. It is implemented using standard web technologies (i.e. HTML, CSS, JavaScript, WebSockets). This creates a computational representation of the mission that supports distributed, synchronous participation among the people responsible for planning the mission. The human planners can be in different parts of the building, area, region, or world and see updates to their plan in real-time through the web UI. It also means that software agents can monitor the process (e.g., observe inputs and track progress), make recommendations to the human planners (e.g., resolve flight deconfliction issues, suggest routes), and act on their own (e.g., provide threat updates, adjust taskings). In the coming sections, we dive deeper into the design and functionality of Metis.

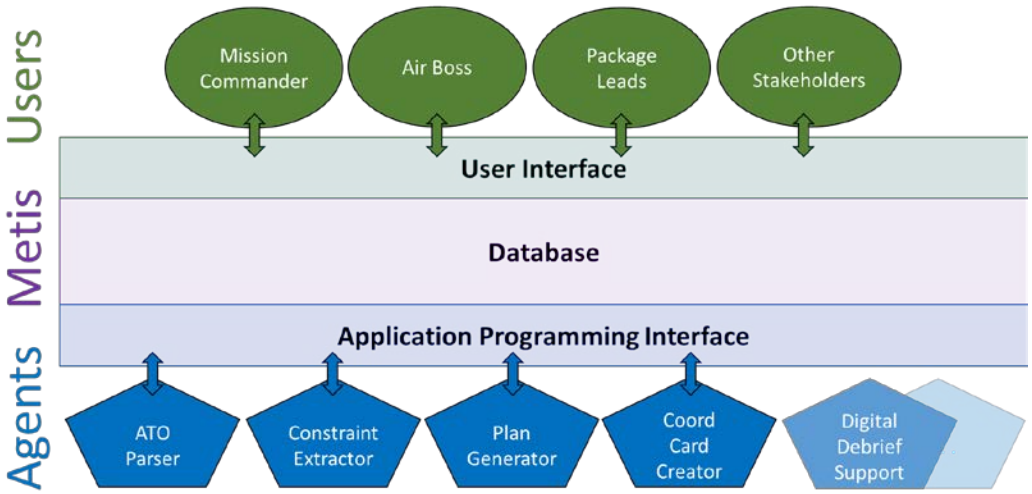

Figure 2. General system structure and relationships among functional components

User Interface (UI)

The UI implementation is written in Javascript, HTML, and CSS. It is designed for flexible use of the system in a variety of hardware options, from handheld mobile devices, to desktop computing, to touch capacitive interactive displays that preserve the form factor of today's mission planning cells (Figure 3).

After the user authenticates into the system, it displays the dashboard menu. The dashboard consists of a set of icons representing the views. Hovering over these icons displays associated tooltips. The Menu icon is a drop-down list that enables a user to export the Coord Card, view the software build version, create a Note template, or logout of the application.

We scope our current implementation by the assumption that an Air Tasking Order (ATO) has already been released. The ATO is the formal means through which a mission's intent, targets, available assets, and timeline are identified. We make the assumption that an ATO has been dropped to a network location. Users browse to a destination folder and upload an ATO to the ATO List and/or delete an ATO from the list. A progress bar is displayed while the ATO is uploaded. The name of the selected ATO is displayed. If the user decides to work on a specific package (PKG) on the ATO, there is a package dropdown list from which users can select. The package might or might not have a Vulnerability (VUL) time displayed next to the VUL. The Aircraft header (AC) displays call signs of all the aircraft in the selected ATO and package.

Figure 3. Next generation mission planning cell hardware configuration. This RDT&E testbed is active in AFRL's Warfighter Readiness Research Division (711HPW/RHA).

The Map view is displayed (Figure 4), by default, when users choose an ATO to process. It enables geospatial 3D visualization of the scenario. An open-source Javascript library called CesiumJS renders the 3D globe and map for visualizing high-resolution global terrain with 3D tiles. Either multi-finger movements on the touch display or buttons on the mouse are used to pan, rotate, and zoom the map. Scalable 3D models of relevant aircraft facilitate interpretability. The Map view displays a dashboard that consists of icons enabling users to toggle between displaying the mission plan in 2D or 3D views. Users may add Markers, Ground Targets, Routes and Waypoints to desired Military Operations Areas (MOAs). The Map offers mission preview using Cesium's built-in timeline and animation widgets. Users can add pre-configured Surface-to-Air Missiles (SAMs) on the map with or without domes that represent a 3D visualization of detection ranges.[1] Users can measure distances and track positions by dragging the mouse across the map. There is a capability to display the format of the coordinates in decimal degrees, unsigned degrees or Military Grid Reference System (MGRS).

Figure 4: Example 2D and 3D views of current mission plan.

The Graphical Air Tasking (GAT) Flow Matrix (GFM), sometimes also now referred to as a Coord Board, displays the parsed ATO in an interactive table (see Figure 5). Data from the raw ATO are parsed and formatted into rows by the ATO parsing agent described later. The dashboard displays icons for sorting the columns, toggling mission color

[1] All SAM site detection ranges are notional and user configurable in the current unclassified implementation.

on or off, and generating a plan based on the data. Columns in the table can be sorted by different values, the default being Mission Type. The rows and columns of the table can be moved and repositioned by a simple drag and drop mouse action. Certain cells are editable and edited cells display an orange mark in the upper right corner until approved by an authorized person, such as a mission commander, air boss, or designated board keeper. Hovering over the edited cell displays a tool tip with information about the last edit.

To enable real-time data sharing and communication, three more views are introduced. They are templates for taking notes, answering intel. questions, and chatting with distributed teammates. Pre-designed Note templates can be used for each ATO and/or at the package level. The Note templates can be further customized. Edited Notes are independent objects in the database to prevent overwriting.

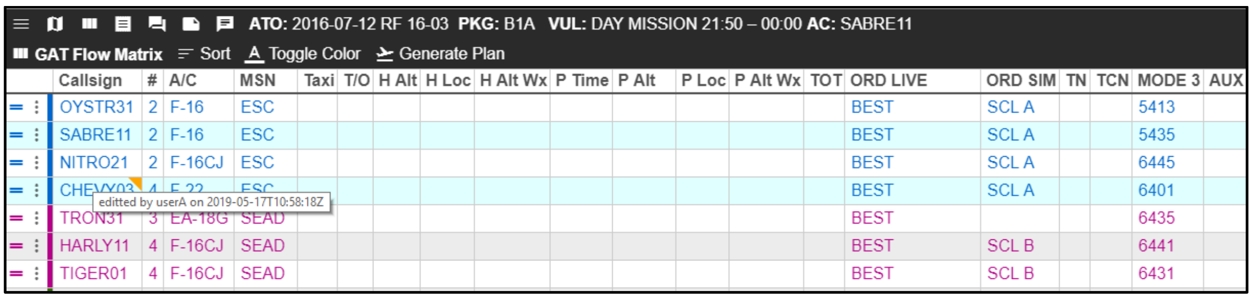

Callsign	#	A/C	MSN	Taxi	T/O	H Alt	H Loc	H Alt Wx	P Time	P Alt	P Loc	P Alt Wx	TOT	ORD LIVE	ORD SIM	TN	TCN	MODE 3	AUX
OYSTR31	2	F-16	ESC											BEST	SCL A			5413	
SABRE11	2	F-16	ESC											BEST	SCL A			5435	
NITRO21	2	F-16CJ	ESC											BEST	SCL A			6445	
CHEVY03	4	F-22	ESC											BEST	SCL A			6401	
TRON31	3	EA-18G	SEAD											BEST				6435	
HARLY11	4	F-16CJ	SEAD											BEST	SCL B			6441	
TIGER01	4	F-16CJ	SEAD											BEST	SCL B			6431	

Figure 5: GFM view displaying an ATO table, package, VUL and aircraft, with colors turned on and the tooltip showing information about the edited cell.

Application Programming Interface (API)

Considerations of reliability, scalability and adaptability led to the choices of our specific Metis software components. After consulting with likely near-term technology off-ramp customers and end user operational subject matter experts, it was clear we required a web-based system as opposed to a desktop solution, due to the need for a distributed, real-time system bereft of the setup headaches associated with a desktop setup option. We also recognized that when a new agent becomes available we require the ability to integrate the new functionality in an efficient and timely manner. This meant we needed to minimize the setup time with potential collaborators in extending cross-domain functionality.

For client/server communication we exploit the WebSocket communications protocol. WebSocket is a modern approach to dealing with the complexity of sending information between the front-end and back-end of a software system, as well as between Metis and other systems, and creates a seamless experience for the front-end user. The client computer first makes a request to initiate communication with the server. After contact is made, communication switches to a two-way protocol. We employ a subscription-based model for WebSocket. This means the client only receives notices and data pertinent to the client's current subscription, cutting down on unnecessary internet traffic. For instance, the client first requests a WebSocket handshake (i.e., subscription) for a particular ATO (or mission). Once this initial communication is established, the server then knows that the client only needs data and updates for this ATO. The upshot of this setup is that each of the clients in the pool of clients subscribed to that ATO receives real-time updates in the planning stages of the mission, regardless of each client's locale. The reduced internet traffic provides for excellent client/server response and turnaround times. Complex mission planning generates lots of data. A significant limiting factor in the current process, which involves whiteboards and no centralized data storage, is that it is impossible to achieve real-time updates. Much of the planning process entails face-to-face interaction and “sneakernet” delivery of hardcopy materials to verify information or to gain approval for mission plan edits.

Database

Metis provides a web interface for people, an API for agents, and a database backend to serve both in real-time. For our backend database, we employ MySQL. MySQL is an open-source, state-of-the-art, powerful Relational Database Management System (RDBMS), for use with web-hosted distributed applications, and as mentioned, serves up data in real-time.

MySQL provides for the ability to manage Tables, Views, Stored Procedures, and Functions. Raw data are stored in the Tables. Various presentations of the data are stored as Views. The data may be gathered from multiple tables and sorted or filtered at the user's discretion. Stored Procedures provide for the ability to query the database on specific subsets of the data by sending parameterized query requests to the database. Finally, there is the ability to write complex Functions in order to create data representations derived from raw data in the Tables. Furthermore, MySQL has the ability to define indexes and primary keys for use with data relationships to ensure data integrity, for instance, cascading deletes or updates. All content parsed from the ATO and all content added by human or agent input to the GFM and Map displays are represented in the MySQL Database.

The general design of the system is that Metis provides a service-oriented architecture within which specific agent services can operate. We have a preliminary set of agents now running, as described next.

AGENTS

Currently implemented agents include an air tasking order parser, a reading agent that extracts constraints from mission planning guidance documents, a mission plan generation and validation agent, a planning product development agent, and an automated debrief focus point identification agent. The following subsections describe each of these in more detail.

ATO Parser

This agent ingests the structured data represented in the ATO and parses it into a set of hierarchical objects. These objects and their metadata are compared to data currently stored in the database. If it is found that the ATO already exists in the database, the user has the option to update the existing ATO, overwrite the existing ATO, or to set it up as a new ATO. Depending on the user's choice or if it is a new ATO to the database, the objects then get stored into the MySQL database. The full stack of parsed data is then sent back to the user that uploaded the ATO and metadata are sent to all other users letting them know a new ATO has been added to the database or to update their information about an existing ATO.

Constraint Extractor

To support a more complete digitization of the mission planning process we needed to go beyond merely ingesting the ATO to also support the extraction of mission planning constraints from natural language text documents. Examples of these documents include Special Instructions (SPINS) and the Air Control Order (ACO). TiER 1 Performance Solutions developed this agent (Warwick, Veksler, Chik, & Rodgers, 2018). It is not a complete solution to the natural language challenge in artificial intelligence, but it does deliver a useful prototype constraint extraction functionality built on the Double R theory and model of natural language processing (Ball, 2004) and on prior Air Force investments in chat-based communications capabilities for synthetic teammates (Ball, Myers, Heiberg, Cooke, Matessa, Freiman, & Rodgers, 2010; McNeese, Demir, Cooke, & Myers, 2018; Myers, Ball, Cooke, Freiman, Caisse, Rodgers, Demir, & McNeese, 2018). The agent searches specific directories to see if the planners have uploaded any new supporting documents. When a document is found it reads the document and identifies sentences that are critical to the mission planning process. The agent uses these sentences to create mission planning constraints. It sends these constraints to the MySQL database, where they are available to guide any processes supporting the mission planning process.

Plan Generator

With an assortment of relevant knowledge already pulled from the ATO and other supporting documents, it is possible to begin the planning process. This is where things begin to get more complicated and reliant on comprehension and creativity. Anything we can do to accelerate this process and improve planning outcomes can be a win for future operations. In that spirit, we have a prototype agent implementation capable of supporting plan generation at the level of resolution available on the GFM and Map displays. The initial baseline code for this agent was also implemented by TiER 1 Performance Solutions (Warwick et al., 2018), and we have begun extending that agent's capabilities. It has two core functions: (1) creating an initial mission plan and (2) checking the validity of any components of the plan input by the human mission planners. The agent works backwards from target locations and strike times to generate push locations and altitudes, and all the way back to the precise times individual planes should taxi for takeoff. Although this plan may not be the final solution, it completes the majority of the "administrative" work for planning a mission. This gives planners more time to use their human creativity to focus on innovative, integrative tactics, rather than mundane (but necessary) details like taxi and takeoff times. The agent maintains a representation of the last approved version of the plan and takes in each change to the plan one at a time. It modifies the plan and then runs a series of algorithms to verify that there are no issues with the new version, such as violations of 3-way de-confliction or other mission policies and guidance. If an issue is found, the agent identifies it in the interface, along with a possible solution to fix the issue. However if there are no issues it sends positive verification and the new plan to the database.

Coord Card Creator

At the end of the mission planning process, the Coordination Card (i.e., Coord Card) is generated. This contains all relevant information for the pilots and serves as an onboard reference during mission execution. The current standard is that a couple of Airmen manually generate the Coord Card. We implemented an agent that takes all approved mission plan data and drops it into an Excel file ready for printing, i.e., the Coord Card. This saves time and decreases opportunities for the introduction of human error. The agent also gives the user an option to print the card in the event of a network outage, which helps to preserve current information and planning progress.

Digital Debrief Support

Most of our emphasis so far, both in the research and development program and in this paper, has been on innovative new capability for the planning portion of the mission cycle. However, with the progress described above on the implementation of Metis and some preliminary intelligent agents, our focus has expanded to include the debrief. The value of a good debrief is well-known and non-controversial. Therefore, there are various existing programs committed to using mission execution data to support digital debrief. Our goal is to leverage those while addressing important niche gaps in connecting a digital representation of the mission plan to those debrief infrastructure technologies. The warfighter is seeking and we are searching for the answer to how we more effectively execute the mission, and in that context we are investing in R&D to accelerate and improve the use of mission plan and mission execution databases to extract debrief focus points (DFPs) and lessons learned. In a recent partnership with a team comprised of researchers from Lockheed Martin (LM) and Massachusetts Institute of Technology (MIT), they developed and delivered an approach that uses Linear Temporal Logic to represent a mission plan (Craven, Oden, Landers, Shah, & Shah, 2018, 2019; Oden, Craven, Landers, & Macannuco, 2019). This representation and data from simulated mission execution are input to an agent that uses a combination of recurrent neural network (RNN) classification and Bayesian specification inference to identify mission phase, aircraft state, and objective completion as interpretable data for debrief.

FORMATIVE EVALUATION

Our research and development for mission planning and debrief is ongoing. At this point we have a subset of the planned system implemented as prototype capabilities. It is too early for statements about end-state benefits and summary accomplishments. However, all publications are progress reports, and in today's more agile research, development, test, and evaluation (RDT&E) environments the focus has shifted to iterative, formative evaluations. Here we describe the general metrics in which we are interested and the currently available results.

In general, as noted above, value to the warfighter is measured in terms of improvements to the efficiency and effectiveness of mission planning and debrief. To begin evaluating our progress toward these benefits we took our system out to Nellis AFB for evaluation at Red Flag 19-1. Red Flag is the USAF's premier large force exercise (LFE) training event, involving dozens of weapon systems and hundreds of personnel. As this was the first formative evaluation of our technology, and the importance of a successful exercise is of the utmost importance, we employed a subject matter expert (SME) with prior Red Flag mission planning experience to "shadow plan" along with the real exercise. Our SME was given a short introduction to Metis and then planned that day's mission using the system. Thus, this evaluation was completed on a non-interference basis with current Red Flag exercise processes, but was done in parallel with their activities. A key metric for comparing the two planning methods (i.e. traditional vs Metis) is the time to the first complete draft plan. Testing included two mission planning cycles. The first mission was Offensive Counter Air (OCA), and the second was Defensive Counter Air (DCA). In both cases, planning was faster in our digital infrastructure, with average savings of 130 minutes. Between-subjects designs of this sort have their inherent weaknesses, but some of the current implementation provides clear advantages over traditional processes. For example, mission commanders at Red Flag spend approximately 90 minutes printing, reading, and manually transcribing content from a hardcopy ATO to the whiteboards on the wall in the Mission Planning Cell. Our ATO parsing agent, by comparison, accomplishes that goal in just a few seconds.

Other evaluations have been accomplished where possible during development of the component capabilities described above. For instance, an evaluation of the Plan Generator agent suggests that most of its output to the Coord Board is as similar to SME-generated content as the SMEs are to each other (Warwick et al., 2018). This is a positive outcome, given that there are many possible plans that are sufficiently safe to execute and approximately equally likely to be successful, and given that it is impossible to prove optimality in these stochastic, uncertain mission environments. The Digital Debrief Support agent provides another evaluation example. Results show 0.92 to 0.96 accuracy in correctly classifying mission phases and aircraft states based on simulated large-force mission data – an important starting point for automated generation of debrief focus points (Craven et al., 2019).

FUTURE RESEARCH AND DEVELOPMENT

As part of our ongoing core applied research, the Mission Planning and Debrief testbed and current implementation of associated agent technologies provides a foundation on which to create the future. From a research program planning perspective, the concept is to use the testbed both as a R&D platform of its own, as described in this paper, and also as a leverageable infrastructure for new investments. These may be more on the fundamental research end of the continuum or more on the advanced technology end, as time, attention, and funding allow. We end the paper with a brief description of some intended directions.

Interactive Task Learning

People rapidly learn to do new things through interactions with others and with their environment. We both learn and instruct through these interactions. Historically, the artifacts we created have not had any means of participating proactively and intentionally in these learning and teaching processes. They have been tools we can use to support learning and teaching, but not capable of doing so themselves. In recent years some ideas have begun to emerge (Laird et al., 2017) around the ideas of Interactive Task Learning (ITL). This involves breaking that historical mold in the implementation of our machine artifacts, in order to make progress on a vision of the future in which humans, robots, and agents are able to rapidly learn and teach each other entirely new tasks through natural interaction. The most complete treatment of these ideas is available in a book released earlier this year (Gluck & Laird, 2019a). The book is in the Strüngmann Forum Series and is an exploration of challenges, open questions, and implications from multiple scientific disciplines for achieving scientific and technological progress in this direction. The introductory chapter emphasizes the issues of pace, persistence, and partnering that emerged in the Forum discussions and writing of the book, as well as the deep challenge of understanding (Gluck & Laird, 2019b). Enduring challenges such as these can focus new research efforts and writing (e.g., Hough & Gluck, in press). The connection to Mission Planning and Debrief is that the testbed provides a concrete context in which to ground new innovations related to ITL. Wray et al. (2019) describe interactivity, generalization/transfer, sufficient complexity, and domains that "matter" as desirable characteristics in environments for ITL research. Mission Planning and Debriefing are highly interactive tasks, provide ample opportunity for knowledge generalization and transfer, are manageably complex, and definitely matter to our current and future warfighters.

Mission Preview/HoloLens

One of our future goals is to provide synchronized preview of the mission across multiple clients and diverse hardware profiles. This would allow mission planners to see a fly-through of the mission to check for de-conflictions and tactical approach issues. Animated preview is a quick and effective way to achieve this. The preview would be synchronized across all connected users allowing for real-time distributed communication about the mission. Beyond displaying the preview on just the map display, we plan to display it across other technology mediums, such as the Microsoft HoloLens. The HoloLens is an augmented reality device that allows the user to see digital information displayed in the real world. The ability to see both the synthetic and physical environment simultaneously is important in the context of a congested and dynamic environment like a Mission Planning Center, where it might be inadvisable to have people "blinded" by full visual immersion in virtual reality. For the HoloLens we have developed the groundwork for displaying terrain data with animated assets flying over, building on Lockheed Martin's digital SandTable technology (United States Navy Office of Naval Research [ONR], 2017). This allows planners to move around and physically interact with the preview rather than through conventional means such as a mouse and keyboard.

Tech Maturation and Transition

Our preliminary results are encouraging, and we are enthusiastic about the possibility of near-term technology off-ramps that can benefit training at large-force exercises. However, the end goal really is to provide technology options that benefit operations and provide flexible force readiness more broadly. We are now in the early stages of exploratory extensions of the prototype capabilities described in this paper. The goal is to move them toward multi- and all-domain operations and a broader set of functional mission types and associated assets. For example, we are collaborating with others in AFRL in an effort to inject cyber effects into the mission planning process. Our current software implementation and documentation afforded us the ability to provide functionality and guidance to facilitate that integration. As multi-domain operations become an increasing focus of development and research, we anticipate these same tools will mature to handle the combined complexities of space, ISR, and the joint fight.

ACKNOWLEDGMENTS

The R&D described in this paper is funded by the U. S. Air Force Research Laboratory. The views expressed in this paper are those of the authors and do not reflect the official policy or position of the United States Air Force, Department of Defense, or United States Government. This paper is cleared for public release, distribution unlimited, 88ABW-2019-3704.

REFERENCES

Ball, J. (2004). A cognitively plausible model of language comprehension. *Proceedings of the 13th Conference on Behavior Representation in Modeling and Simulation*, (pp. 305-316). ISBN: 1-930638-35-3

Ball, J., Myers, C. W., Heiberg, A., Cooke, N. J., Matessa, M., Freiman, M., & Rodgers, S. (2010). The synthetic teammate project. *Computational & Mathematical Organization Theory, 16*(3), 271-299.

Craven, P. L., Oden, K. B., Landers, K. J., Shah, A. J., & Shah, J. A. (2018, November). *Man-Machine Interoperation in Training for Offensive Counter Air Missions*. Paper presented at the Interservice/Industry Training, Simulation and Education Conference 2018, Orlando FL.

Craven, P. L., Oden, K. B., Landers, K. J., Shah, A. J., & Shah, J. A. (2019, December). *Man-Machine Interoperation in Training for Large Force Exercise Air Missions*. Paper presented at the Interservice/Industry Training, Simulation and Education Conference 2019, Orlando FL.

Department of Defense. (2018, January 19) *Summary of the 2018 National Defense Strategy of the United States.* Retrieved from https://dod.defense.gov/Portals/1/Documents/pubs/2018-National-Defense-Strategy-Summary.pdf

Gluck, K. A., & Laird, J. E. (Eds.) (2019a). *Interactive task learning*. Strüngmann Forum Report, vol. 26, J. Lupp, series ed. Cambridge, MA: MIT Press.

Gluck, K. A., & Laird, J. E. (2019b). Looking forward to interactive task learning. In K. A. Gluck & J. E. Laird (Eds.), *Interactive task learning*. (pp. 1-7). Strüngmann Forum Report, vol. 26, J. Lupp, series ed. Cambridge, MA: MIT Press.

Hough, A., & Gluck, K. A. (in press). The understanding problem in cognitive science. *Advances in Cognitive Systems.*

Laird, J. E., Gluck, K. A., Anderson, J. R., Forbus, K., Jenkins, O., Lebiere, C., Salvucci, D., Scheutz, M., Thomaz, A., Trafton, G., Wray, R., Mohan, S., & Kirk, J. (2017). Interactive task learning. *IEEE Intelligent Systems, 32*(4), 6-21, (invited).

McNeese, N. J., Demir, M., Cooke, N. J., & Myers, C. W. (2018). Teaming with a synthetic teammate: Insights into human-autonomy teaming. *Human Factors*, *60(2)*, pp 262-273. DOI: 10.1177/0018720817743223

Myers, C.W., Ball, J., Cooke, N.J., Freiman, M., Caisse, M., Rodgers, S., Demir, M., & McNeese, N. (2018). Autonomous intelligent agents for training: Making the case for synthetic teammates. *IEEE Intelligent Systems.* DOI: 10.1109/MIS.2018.2886670

Oden, K., Craven, P., Landers, K., & Macannuco, D. (2019). *Learning for Man-Machine Interoperation and Training* (Report AFRL-RH-WP-TR-2019-0023). Wright-Patterson Air Force Base, OH: AFRL.

United States Air Force. (2019) *United States Air Force 2030 Science and Technology Strategy.* Retrieved from https://www.af.mil/Portals/1/documents/2019 SAF story attachments/Air Force Science and Technology Strategy.pdf

United States Navy Office of Naval Research. (2017). *HoloLens SandTable User Manual* (Version 2.9 for SandTable 3.9). Orlando FL: Lockheed Martin Rotary & Mission Systems.

Warwick, W., Veksler, B., Chik, K., & Rodgers, S. (2018). *SynCoor: Synthetic Coordination System for Airman-Autonomy Mission Planning and After-Action Review* (Report AFRL-RH-WP-TR-2018-0071). Wright-Patterson Air Force Base, OH: AFRL.

Wray III, R. E., Taatgen, N. A., Lebiere, C., Pastra, K., Pirolli, P., Rosenbloom, P. S., Scheutz, M., Stewart, T. C., & Wiles, J. (2019). Looking forward to interactive task learning. In K. A. Gluck & J. E. Laird (Eds.), *Interactive task learning.* (pp. 19-51). Strüngmann Forum Report, vol. 26, J. Lupp, series ed. Cambridge, MA: MIT Press.