

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/366393913>

Virtual reality framework for multi-human multi-agent adaptive teamwork

Conference Paper · December 2022

CITATIONS

0

READS

532

5 authors, including:



[Joey Salisbury](#)

Riverside Research

45 PUBLICATIONS 830 CITATIONS

SEE PROFILE



[Ross L. Bobb](#)

Riverside Research

3 PUBLICATIONS 2 CITATIONS

SEE PROFILE



[Virgil Barnard](#)

Riverside Research

12 PUBLICATIONS 4 CITATIONS

SEE PROFILE



[William Casebeer](#)

Beyond Conflict International

45 PUBLICATIONS 726 CITATIONS

SEE PROFILE

Virtual reality framework for multi-human multi-agent adaptive teamwork

Joseph P. Salisbury
Riverside Research
Lexington, MA
jsalisbury@riversideresearch.org

Ross L. Bobb
Riverside Research
Beavercreek, OH
rbobb@riversideresearch.org

Virgil O. Barnard IV
Riverside Research
Beavercreek, OH
vbarnard@riversideresearch.org

William D. Casebeer
Riverside Research
Lexington, MA
wcasebeer@riversideresearch.org

David M. Huberdeau
Riverside Research
Lexington, MA
dhuberdeau@riversideresearch.org

ABSTRACT

Individualized, adaptive intelligent technologies that continuously promote the emergence of team cohesion in novel groups of humans and intelligent agents are necessary to maximize collective adaptation to rapidly evolving environmental demands. Here, we describe an iterative three-phase human-centered design and evaluation framework to aid in the development of teamwork-promoting autonomous capabilities. Qualitative and quantitative feedback from each phase can be used to improve the fidelity of simulations for training, autonomy performance, and the design of user interfaces. In the initial phase, synthetic agents are used to approximate human behaviors and provide initial training of AI models. In the second phase, individual human operators are introduced to evaluate the capability in a simulated operational environment. In the third phase, networked virtual reality clients allow human teams to collaborate on tasks in an immersive, physics-based simulation environment to evaluate how the capability may complement and enhance teamwork performance. To demonstrate this framework, we developed a virtual reality combat simulation demonstrating a rifle-mounted fire control system with target detection and tracking algorithms. Test user and stakeholder feedback was collected and reviewed to establish requirements for an intelligent decision-making aid to fuse data from individual fire control systems and coordinate target allocation and threat prioritization. Miniature unmanned aerial vehicles were incorporated to assist target tracking. Combat simulations are used to train algorithms to detect and track threats, predict outcomes, and provide feedback to coordinate squad tactics based on individual and group factors. Integration of these artificial intelligence capabilities into the virtual reality environment enable human-in-the-loop evaluation of their impact on multi-human multi-agent teamwork. By providing a continuous environment for model training and human-in-the-loop evaluation in virtual reality, teamwork autonomies can be agilely developed centered around human factors to improve both performance and teammate acceptance.

ABOUT THE AUTHORS

Dr. Joseph P. Salisbury is an AR/VR Research Scientist at Riverside Research. He received a Ph.D. in Neuroscience from Brandeis University in 2013 and has since focused on AR/VR, human factors, and physiological monitoring.

Mr. Ross L. Bobb is an AR/VR Applied Scientist at Riverside Research. He has a master's degree in Electro-Optics from the University of Dayton. His work includes synthetic data generation and VR for performance assessment.

Mr. Virgil O. Barnard IV is a Senior Machine Learning Scientist at Riverside Research. He is completing his Ph.D. in Computer Science from the University of Kentucky. A graduate of the U.S. Marine Corps Scout Sniper School, he works on object identification and activity classification problems across multiple sensor modalities.

Dr. William D. Casebeer is Director of Artificial Intelligence and Machine Learning at Riverside Research. Dr. Casebeer has decades of experience leading teams to develop solutions to pressing national security problems.

Dr. David M. Huberdeau is an AI/ML Engineer at Riverside Research. He received a Ph.D. in Biomedical Engineering from Johns Hopkins University and completed post-doctoral work at Yale University. Dr. Huberdeau's current research includes human behavior modeling and reinforcement learning.

Virtual reality framework for multi-human multi-agent adaptive teamwork

Joseph P. Salisbury
Riverside Research
Lexington, MA
jsalisbury@riversideresearch.org

Ross L. Bobb
Riverside Research
Beavercreek, OH
rbobb@riversideresearch.org

Virgil O. Barnard IV
Riverside Research
Beavercreek, OH
vbarnard@riversideresearch.org

William D. Casebeer
Riverside Research
Lexington, MA
wcasebeer@riversideresearch.org

David M. Huberdeau
Riverside Research
Lexington, MA
dhuberdeau@riversideresearch.org

INTRODUCTION

Ongoing advances in artificial intelligence and machine learning (AI/ML) are enabling autonomous agents to participate as collaborative teammates that contribute to teamwork effectiveness and performance (Larson & DeChurch, 2020; O'Neill et al., 2020; Seeber et al., 2020). Autonomous agents (or “autonomies”) can be leveraged to support teamwork activities including team coordination, task reallocation, and continuous interactions between humans and other autonomous agents (Frame et al., 2020; Madni & Madni, 2018; Roth et al., 2019). Teamwork can be viewed as a multilevel process that includes individual taskwork as well as individual- and team-level states (e.g., cohesion, shared mental models, shared situation awareness) and processes (e.g., coordination and communication) that influence team performance and effectiveness (DeCostanza et al., 2018; Kozłowski & Klein, 2000; Marks et al., 2001; Salas et al., 2007). An autonomy that can sense, process, and respond to individual- and team-level states can use this to adapt to evolving team dynamics and provide personalized and complementary assistance to enhance both individual and team-level performance (Grimm et al., 2018; Mait et al., 2017; Schaefer et al., 2021).

Autonomies that are both adaptive and robust are especially critical for effective team performance when operational contexts are complex, austere, and involve high, potentially life-threatening, risks (Schaefer et al., 2021). When autonomies operate in the same environment as humans, there is risk of conflict (e.g., the operator and the autonomous component work at cross-purposes), discoordination (e.g., the autonomous component failing to account for operator actions), and cognitive bias in operator decision making (e.g., the operator choosing sub-optimal actions based on overestimating the likelihood of rare events). To address this, human-aware planning requires an autonomy to have a representation of the perceived mental model of humans collaborators – an artificial “theory of mind” (Rabinowitz et al., 2018; Williams et al., 2022). Theory of mind could be utilized to infer operator intent as well as provide explainable plans that human operators can comprehend and trust (Chakraborti et al., 2018). When interacting with multiple human teammates, the autonomy must be able to identify an explainable policy – a strategy in pursuit of a goal – that is consistent across the mental models of all teammates to avoid confusion and loss of trust. Individual augmented reality (AR) displays can provide personalized feedback to address this problem (Sengupta et al., 2018).

One domain where autonomous agents could provide adaptive and individualized teamwork enhancements via AR displays is the dismounted rifle squad. Intelligent decision-making aids (IDAs) implemented in dismounted systems could be used to identify and recognize threats, predict adversarial force position and movement, and highlight ‘danger areas’ (Geuss et al., 2019). This information could be displayed using AR on a weapon-mounted optic, such as the next-generation squad weapon fire control (NSGW-FC, Figure 1A). The NSGW-FC includes a variable magnification optic, backup etched reticle, laser rangefinder, ballistic calculator, atmospheric sensor suite, compass, Intra-Soldier Wireless, visible and infrared aiming lasers, and a digital display overlay. Using the NSGW-FC with additional sensors enables an aided target recognition (AiTR) system that can detect and highlight threats, prioritizing operator attention. Fusing data from multiple sources, including squad embedded autonomies such as the Black Hornet Personal Reconnaissance System (PRS, Figure 1B), an IDA could coordinate target tracking within and across squads, direct squad movement, formation, and fields of fire, and facilitate target prioritization and allocation. IDAs with access to the roles, capabilities, limitations, behaviors, and physiology of their human teammates could leverage this data for more optimal decision making. For example, behavioral data can be used to anticipate actions and needs and train

small unmanned aerial or ground systems (sUAS/sUGS) to support the dismounted squad (Lance et al., 2020). IDAs aware of the limitations of AI/ML capabilities could help human teammates calibrate the level of trust they should give to information provided (Scielzo et al., 2021). Previously, we developed a virtual reality (VR) testbed to address several needs critical to an effective dismounted soldier AiTR-AR system (Bobb et al., 2022). Our main objectives were to provide a platform for:

1. Generating operationally relevant synthetic data for training AiTR algorithms.
2. Rapid prototyping of AiTR-AR capabilities and test case scenarios.
3. Objective evaluation of how the system impacts the operator under simulated operational conditions.

Each of these capabilities contribute to our ability to design, evaluate, and improve upon the dismounted soldier AiTR-AR system with an iterative human-centered approach. For example, while AiTR-AR displays are intended to help optimize attention resources, there are numerous ways they may fail to adequately consider the cognitive mechanisms underlying allocation of attentional resources and adversely impact visual cognition and search (G. Larkin et al., 2020; G. B. Larkin et al., 2020). Likewise, the effectiveness of AiTR algorithms can be confounded by a variety of parameters introduced by challenging operational conditions, such as environmental obscurants and hardware limitations (e.g., sensor accuracy/resolution, processor speed, size, weight, power consumption, etc.). Therefore, it is necessary to train AiTR algorithms on data representative of an operator's perspective and validate early in the design process they are performant given operational constraints. Through this operator-centric approach, our VR testbed enables continuous improvement in each sub-system that impact human-machine synergies and performance gains.

To demonstrate how the VR testbed could be used to objectively evaluate AiTR-AR designs, we completed a proof-of-concept system evaluation that measured operator physiology and task performance with and without AiTR-AR in a combat scenario (Bobb et al., 2022). Evaluation participants exhibited reduced stress and improved performance with AiTR-AR. Participants reported the system aided their ability to detect, track, and engage threats. Having demonstrated how VR could be used to evaluate a dismounted soldier AiTR-AR, we were motivated by stakeholders to explore capabilities that enhanced human-machine teaming synergies at both the individual- and squad-level. In this paper, we describe initial progress toward developing a revised VR testbed for the design and evaluation of IDAs for the dismounted soldier squad. This teamwork coordination IDA would provide adaptive AiTR capabilities that prioritize and allocate targets based on individual- and team-level factors. To achieve this, we:

1. Refined system user needs based on feedback from system evaluation participants and stakeholders
2. Present a three-phase framework utilizing VR to iteratively design and evaluate a teamwork coordination agent, and describe progress on using this framework to develop an IDA to coordinate squad actions

We describe how the proposed development framework can be applied to design and evaluate multi-human multi-agent AI/ML capabilities, using an IDA for rifle fire team coordination as example.

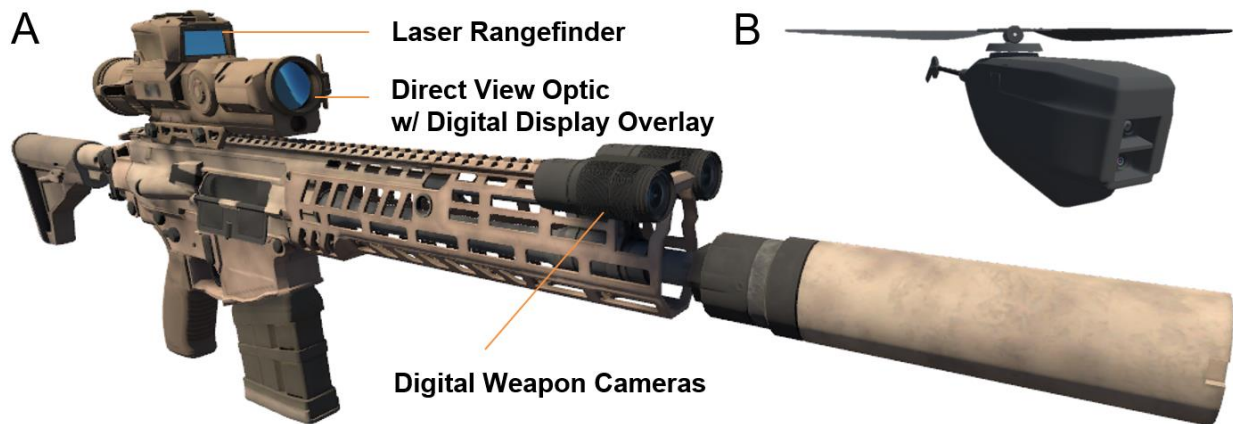


Figure 1. Technologies used in demonstration of multi-human multi-agent adaptive teamwork in VR. A. Next generation squad weapon (Sig Sauer MCX-Spear) and fire control (Vortex Optics XM157) with digital cameras for target tracking. B. The 16.8 cm long Black Hornet UAV can provide situation awareness for the dismounted squad.

METHODS

VR Environment

The VR scenario was developed in Unity3D (2019.4.13f1) and evaluated using the Oculus Quest 2. A full description of the system and user evaluation can be found in our prior report (Bobb et al., 2022). Briefly, evaluation participants were positioned in a guard tower with (AI, non-human) squad members positioned in guard towers on opposing corners of the facility (Figure 2A). Participants were informed their primary objective was to guard crates from adversarial forces. Participants received training on mission objectives, their weapon, and the AiTR-AR system being evaluated (Figure 2B) from RITA (Real-time Intelligent Training Assistant). Participants completed two rounds of the combat scenario, with and without AiTR system, with each condition lasting five minutes. The order of these conditions was randomized across participants. A chest-worn heart rate monitor (Polar H10) was used to measure heart rate activity during task performance. Participants were trained in a tactical breathing exercise by RITA as a method for baselining stress between combat conditions. Various task-related performance metrics were collected, including the number of shots fired, participant accuracy, number of shots on participant, and number of crates remaining. An outline of the evaluation experience is shown in Figure 2C.

Analysis of Participant Performance and Feedback

After completing the VR experience, participants (N=12) were asked about their most and least favorite aspects of the experience, what they found most challenging, what they found most confusing, and what they would change about the experience. Previously, we analyzed the responses to these prompts around the specific benefits and drawbacks of the AiTR system being evaluated. We revisited participant responses to identify needs that may be addressed by an IDA with squad-level information. Participant responses to each prompt were reviewed and relevant feedback was coded to identify common themes regarding what participants found challenging, confusing, or recommend improving beyond the core AiTR capability being evaluated. Preliminary requirements and concept designs were generated based on participant feedback. We also reexamined task performance metrics to identify areas of improvement.




Figure 2. Summary figure of Bobb et al., 2022 evaluation. A. Layout of combat scenario environment. Blue boxes show friendly guard tower positions relative to the participant's guard tower view (yellow outlined inset). B. Example AiTR-AR highlighting through weapon-mounted optic. C. System evaluation timeline for participants.

Stakeholder Needs Assessment

Semi-structured interviews were conducted with external stakeholders as part of a design review. Interview probes were adapted from the Human-Machine Teaming Knowledge Audit (Dominguez et al., 2020; McDermott et al., 2018). Three external stakeholders were interviewed to refine additional needs for the AiTR system. Stakeholders were specifically interested in developing capabilities for a weapon-mounted AiTR-AR system for dismounted soldiers. Prior to the interview, preliminary requirements and concept designs based on participant feedback were presented to an internal subject matter expert. After the interview with external stakeholders, feedback was reviewed to refine the documented set of use cases and system requirements. Based on these requirements, a plan was specified for how to iteratively develop and evaluate the system using the VR testbed developed for the prior system evaluation.

RESULTS

Participant Needs After VR Evaluation

Previously, we focused our analysis of participant responses on the usability of the AiTR-AR system itself. In reviewing participant responses with the goal of identifying more general challenges encountered, we identified two major themes: 1) difficulty identifying where enemy fire was coming from to maintain adequate cover; and 2) the need for better collaboration with squad members to engage targets out of range. Almost all participants reported difficulty engaging targets that were too distant while avoiding incoming fire. For example, Participant 5 reported this and noted “more help from my teammates in other [guard]towers” was necessary. Participant 3 noted the most challenging aspect of the experience was “locating distant enemies, [and] remembering to take cover.” Even when not being targeted by an adversary, Participant 1 noted he was frustrated he could not effectively engage an enemy he spotted due to obstacles in his line-of-sight -- “I had to wait for the tower guard to take him out.” As there was no method to communicate with team members, this limited the potential for participants to coordinate.

In terms of performance, participants generally failed at the task of defending assets from adversarial forces. Two-thirds of participants lost all assets they were defending on the unaided condition and, even with the AiTR-AR system, half of participants lost all assets. This suggests that while the AiTR-AR system helped participants engage targets more effectively and, thus, better defend the assets, there was still considerable room for improvement. A more effective AiTR-AR system could help prioritize which targets to engage with, reducing cognitive load on the operator and keeping them focused on mission objectives. Similarly, we saw that the AiTR-AR system led to significant reductions in the amount of enemy fire taken by the participant. We hypothesized that the AiTR-AR system enabled participants to identify and engage with targets more quickly during intervals when they may be exposed to enemy fire. A system that helped coordinate targets across the squad and prioritize which targets were necessary for a particular squad member to engage could further alleviate exposure times while preventing adversarial gains. Finally, when comparing participants with and without the AiTR-AR, while we observed significantly more lethal shots with aid of the AiTR-AR system ($p = .03$), we did not see a difference in overall shot accuracy ($p = .82$). However, a system that coordinated squad fire by assigning targets to squad members who were best positioned to engage could help improve accuracy and, in turn, lethality. Factoring individual skill levels into this target assignment could also boost individual and group-level lethality.

Stakeholder Needs

Stakeholders were primarily interested in refining computer vision capabilities for the AiTR-AR system, including improving the robustness of detecting and tracking obscured targets. Stakeholders were also interested in estimating the distance and velocity of tracked targets using weapon-borne passive sensors. Current challenges of the AiTR-AR computer vision capabilities include “jumpy” tracking of targets that should be smoothed to avoid distracting or confusing the operator. Likewise, currently existing tracking algorithms had difficulty with obscuration, which made operators question their utility and reliability. By failing to track partially obscured targets that a human could easily identify (e.g., a person’s leg still being visible behind cover), this failure at basic reasoning was hampering the trust operators had for the technology. Thus, a primary requirement of the system was to ensure detection and tracking of human targets be robust to obscuration. Another challenge was when tracking algorithms failed to persistently track a specific human as they cross paths with other humans. The envisioned system should be able to detect all humans in a scene and then the operator should be able to “tag” one human (i.e., identifying them as a known or observed threat) for persistent tracking. To achieve this, the system should be able to learn the features of that target on the fly.

For hardware, the system will have a wide-angled camera to maintain tracking targets in a wide field-of-view. In the digital overlay, stakeholders envisioned an arrow in the digital overlay could guide operators to the target, pointing them in an appropriate direction. One stakeholder likened it to the challenge of guiding someone through verbal cues to acquire a target down range, such as following the tree line to a certain distance before scanning across the field.

The ability to passively range targets (i.e., without the use of the laser rangefinder) was also of interest, so the ballistic computer could automatically adjust the sight with minimal intervention. In addition to the target range, estimating target heading and velocity would also aid in ballistic adjustments and situation awareness. Providing estimate accuracies to the operator may be helpful to judge whether adjustments can be trusted to increase hit probability.

Stakeholders expressed there was great interest in the AiTR-AR system being able to share data to enable a “common battlefield perspective.” Stakeholders, who were primarily interested in developing capabilities for a weapon-mounted system, said their main priority was to increase soldier lethality. They contrasted this primary objective with the more general goal of increasing situation awareness. For example, an AiTR-AR system that could receive information from elsewhere in the common battlefield network could provide individual dismounted soldiers a targeting queue. By sharing “target reference points,” squad members could maintain tracking targets and “hand off” that target’s location to each other, autonomies (e.g., unmanned aerial/ground vehicles, UAVs/UGVs), or another squad. Stakeholders encouraged us to explore solutions without worrying about the specific capabilities of any given networking solution. When we pressed to understand what a reasonable transmission range could be, we were told to focus on potential performance gains at a variety of ranges, as increases in squad lethality at any scale could help motivate determining an appropriate solution. More importantly, solutions would be constrained by the size, weight, and power limitations that were feasible for individual dismounted soldiers.

Stakeholders agreed the use of simulations, such as virtual reality, were helpful in evaluating first-hand the performance of algorithms. While the VR simulation might not exactly represent the complexities of real-world operations, it is beneficial to have the ability to assess potential performance gains for a capability. This could provide rationale to push development for a field-deployable system.

Revised System Needs and Design

Table 1 shows high-level user needs identified from participants and stakeholders to guide system design henceforth. Utilizing our existing VR testbeds capabilities, we sought an approach to demonstrate how an IDA could aid in coordinating the dismounted squad, as well as squad-embedded autonomies. This team coordination agent could provide a joint plan based on the shared goals of the squad and communicate the joint plan to individual human teammates in an explainable way based on theory of mind techniques to infer mental models of human teammates.

Table 1. High-level user stories for revised system based on participant and stakeholder feedback.

<i>As a dismounted soldier, I want...</i>	<i>...so that...</i>
To evaluate AiTR systems with my squad members	I can understand how they impact group performance
To evaluate an accurate representation of an AiTR in VR	I can learn the new capabilities of the system and determine if it meets my needs for field testing
To evaluate how well the AiTR detects targets	I can ensure it will serve my needs on the battlefield, such as detecting targets under degraded visual conditions
To evaluate how well the AiTR tracks targets	I can ensure it will serve my needs on the battlefield
To evaluate the passive ranging capabilities of the AiTR	I can evaluate how well the AiTR passive ranging capability helps extend my lethality range
To communicate target information with my squad	We can share knowledge and coordinate
To be able to receive objective metrics of how my team performed on the task	I can track how well the AiTR improves performance, as well as how well my team is trained to utilize the AiTR

An Iterative Three-Phase Approach to Developing Team-Aware Planning Agents

To develop a team coordination IDA, we propose a three-phase human-centered design framework for iterative development and evaluation of personalized, adaptive multi-human multi-agent team coordination agents (Figure 3). Following the principles of human-centered design, the core philosophy of this approach is to enable rapid, hands-on prototyping of AI/ML capabilities by maintaining a continuous design, development, and virtual training environment for both humans and autonomies. The three phases include:

- **Phase 1:** All Synthetic Evaluations
- **Phase 2:** Human-in-the-Loop Evaluations
- **Phase 3:** Team-in-the-Loop Evaluations

Using this approach, human- and team-in-the-loop testing and stakeholder feedback can be obtained rapidly based on phase transitions of minimally viable product increments. This approach is iterative, with models and designs receiving progressive validation, while test results and stakeholder feedback can be used to improve simulation fidelity and system capabilities in a previous phase. While the first phase includes fully synthetic proxies for humans, early testing with human research participants in virtual reality enables the collection of physiological and behavioral data to refine how humans perform under the simulated operational conditions.

Henceforth, we will describe each of the three phases in more detail, including what design and research questions can be addressed at each stage, describe off-the-shelf capabilities that can be leveraged to implement each stage, and describe our preliminary approaches for implementing a system for coordinating individual and squad-level work.

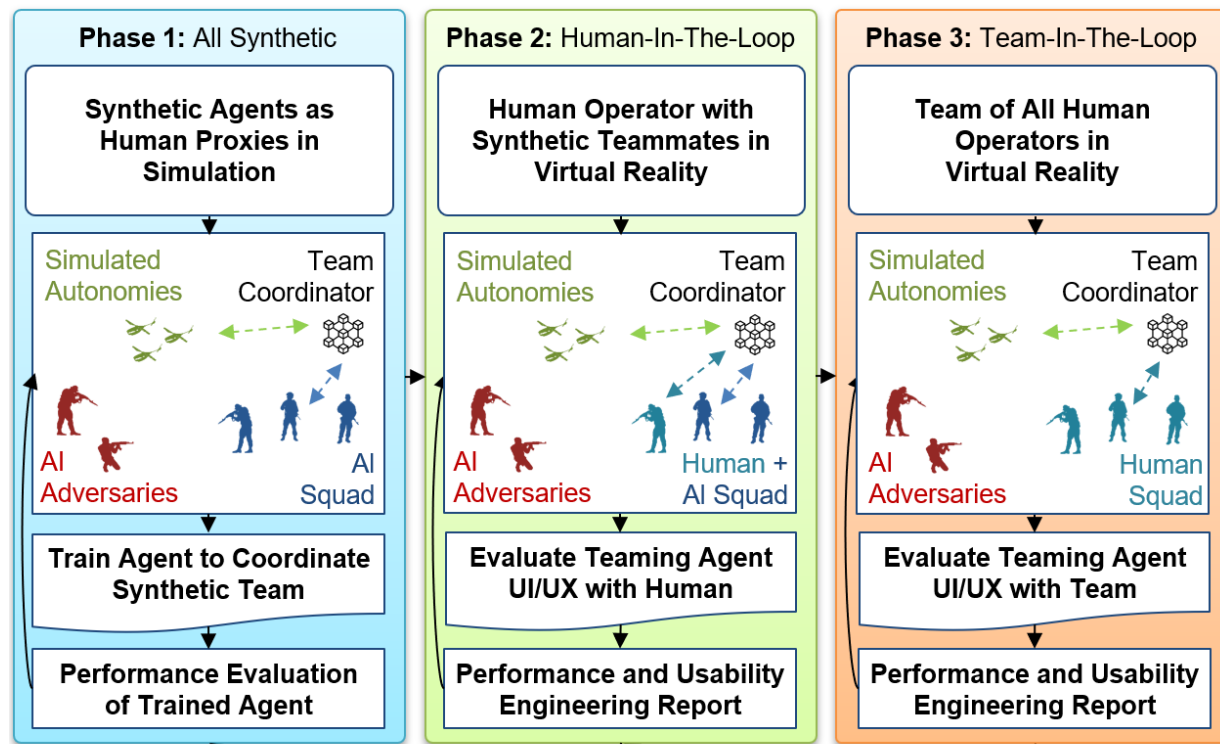


Figure 3. A three-phase iterative approach for human-centered development and evaluation of personalized, adaptive team coordination agents with virtual reality. In Phase 1, synthetic agents are used to approximate human behaviors and train preliminary models for team coordination based on individual states, behaviors, strengths, and limitations. In Phase 2, individual human operators evaluate the UI/UX of the capability in the simulated operational environment. In Phase 3, multiple humans cooperate as a team to evaluate how the capability may complement and enhance teamwork performance.

Phase 1: All Synthetic Evaluations

Phase 1 simulations include synthetic agents, both simulated humans and autonomies, completing a task in an environment (Figure 4). Simulations provide initial training for AI models based on increasingly higher fidelity simulations of agent behaviors and environmental/scenario simulations. Design and development of the simulation enables technologists and stakeholders to refine their common understanding of the task, its requirements, and critical decision points. Importantly, the Phase 1 simulation environment is designed with the intent to provide a simulation environment for human- and team-in-the-loop evaluation. Through this, not only can simulation fidelity and agent quality be validated, but synthetic human behavior can be refined through learning from actual humans in Phases 2 and 3 (e.g., through imitation learning).

In the example of developing a squad coordination agent, observations for the squad coordination agent include the fusion of information from both synthetic squad members and simulated autonomies. The environment is designed with the intention of providing the same operational scenarios for human-in-the-loop evaluation in subsequent phases. Unity ML-Agents is leveraged to provide access between the Unity simulation environment and Python machine learning toolkits (Juliani et al., 2020). Synthetic human characters in the simulation provide data for training a supervisor agent that has access to fused environmental observations. This supervisor agent recommends actions to individual synthetic team members. This supervisor-based hierarchical reinforcement learning approach has been implemented on Unity ML-Agents demo examples and resulted in superior performance results (Cao et al., 2020).

While the Unity ML-Agents framework provides a powerful basis for implementing reinforcement learning paradigms, it includes a wrapper for enabling integration with the OpenAI Gym framework. The OpenAI Gym framework provides a common access point for exposing reinforcement learning components (i.e., state and action space, rewards) to additional frameworks [e.g., OpenAI Gym (Brockman et al., 2016), Stable Baselines3 (Raffin et al., 2021), DeepMind Acme (Hoffman et al., 2020)]. In addition to reinforcement learning packages, other algorithms can be leveraged to refine human-aware planning-based approaches for inferring agent policies (i.e., the strategy that an agent uses in pursuit of goals), implementing predictive analytics of agent behavior (i.e., what will the agent do next? how will the agent respond to change?), and providing explainable joint plans. Finally, trained models can be integrated back into Unity for real-time evaluation via the open neural network exchange (ONNX) format, which can be used to accept models from a variety of ML frameworks including TensorFlow, PyTorch, and Jax.

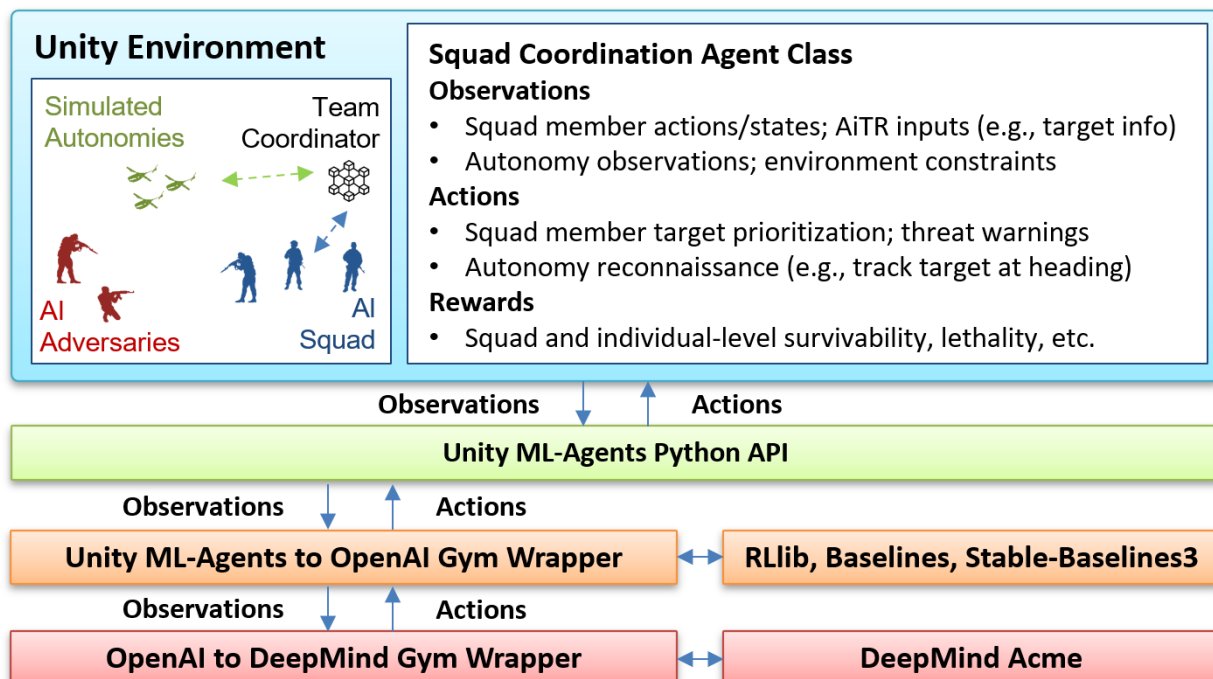


Figure 4. Phase 1 overview. A team coordination agent can be trained to coordinate synthetic squad members and simulated autonomies using a variety of reinforcement learning toolkits.

Phase 2: Human-in-the-Loop Evaluations

As simulations and trained models are validated in Phase 1, model actions must be translated into a user experience and user interface (UI/UX) suitable for Phase 2 human-in-the-loop evaluation (Figure 5). Phase 2 enables both hands-on evaluation of UI/UX designs and AI model performance. Even prior to the development of ML models, UI/UX designs for how an IDA would provide feedback to the operator could be evaluated if heuristic approaches are available. For example, in Bobb et al., 2022, we evaluated an idealized threat detector, as the location of all threats were known in the simulation environment. This allowed us to demonstrate how the impact of the capability may be objectively evaluated, including performance and physiological measures, without the AI capability itself needing to be developed and integrated. We also created a workflow using the Unity Barracuda package (Scott et al., 2021) to integrate functional computer vision algorithms so they may be evaluated in the same environment.

Using this same approach, we can evaluate an idealized squad coordination system. As an initial proof-of-concept, we implemented a scenario with favorable odds for the blue team squad. We limited the number of adversarial forces targeting squad members so that there was always at least one squad member not actively being targeted. This enabled a simple heuristic algorithm to be created, where blue team squad members not under fire could be advised to engage adversaries while others remained under cover. To demonstrate a UI/UX for human-in-the-loop testing, we modified the AiTR-AR interface to guide the operator to targets. We modified the voice AI assistant, RITA, to provide threat level warnings to assist in determining the risk associated with engaging.

With this minimum viable demonstration of the team coordinator, the system can begin evaluation with humans-in-the-loop to further refine and validate the needs of the capability. Hands-on evaluation of the UI/UX and AI can aid in prioritizing the need for additional feedback, such as changes to mission objectives, the status of autonomous capabilities, and personalizing feedback based on operator state estimation, including physiology and behavior. For example, soldier-borne sensors can be used to assess cognitive load and stress to adaptively adjust task load on squad members or minimize distractions in their AiTR-AR system. Weapon-borne sensors can be used to measure operator state, such as resting, tracking a target, transitioning between potential targets, and firing (Lee et al., 2022). Autonomies, such as miniature UAVs, can be directed to targets and provide operators advanced warning regarding target movement patterns behind cover. Shot accuracy measurements taken from the AiTR-AR system could also be used to adjust target assignment within the squad, providing squad members with targets they are best able to engage based on real-time assessment of their performance.

Finally, by integrating human research participants in Phase 2, ML approaches such as imitation learning can be leveraged to train more realistic Phase 1 synthetic agents. More simply, reaction time measures gathered during Phase 2 evaluations can be used to create more realistic Phase 1 agent simulations. These additional constraints during Phase 1 agent training could improve the ecological validity of AI models trained in simulations.

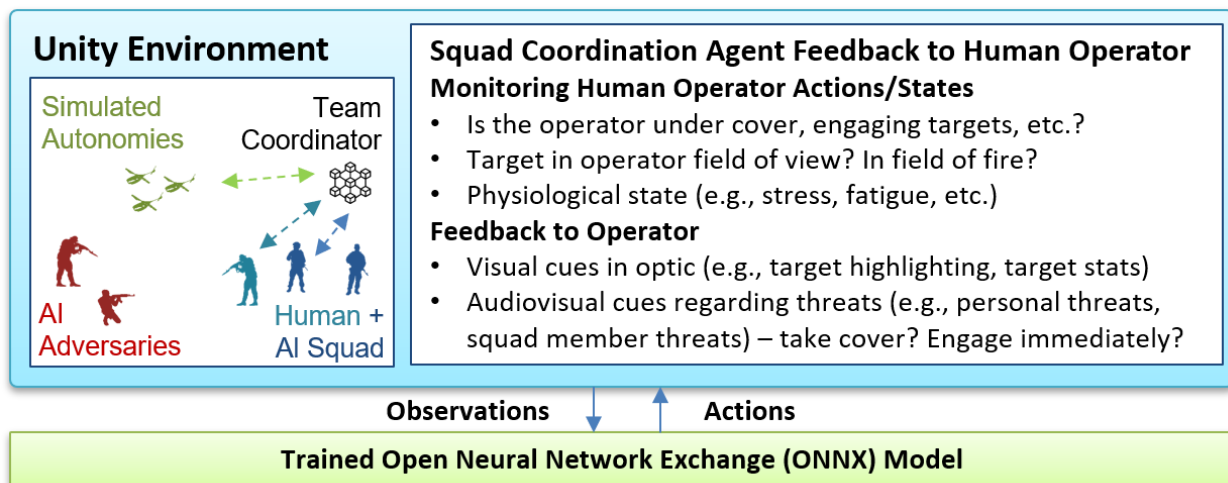


Figure 5. Phase 2 overview. An individual human operator evaluates a team coordination agent coordinating a multi-human multi-agent team to perform tasks in a simulated operational scenario.

Phase 3: Team-in-the-Loop Evaluations

For a team coordination agent to be validated under simulated operational conditions, team-in-the-loop evaluation is necessary to demonstrate all human factors are addressed. Thus, in Phase 3, networking of VR clients with real-time communication channels is used to evaluate performance (Figure 6).

Like Phase 2, Phase 3 team-in-the-loop evaluations can begin using a heuristic approach that does not require model training, only a minimal viable UI/UX to demonstrate how the capability could function under a set of controlled scenarios. From there, a primary objective can be validation of the minimum viable capability. Individual and group feedback can be obtained to refine the system's requirements, as well as the requirements of the evaluation scenario. While it is reasonable to expect a team coordination agent could perform target allocation under some set of assumptions with efficiency exceeding team performance without the agent, there are an endless number of unpredictable scenarios requiring human individual and collective decision-making to adequately assess. Thus, the team coordination agent should be thought of as increasing squad efficiency in a *complementary* way, recommending explainable decisions while also allowing for team communication to override decisions that human judgement can determine is suboptimal (Bansal et al., 2021). Sensor streams, such as voice communication between team members and behavioral actions, can be analyzed to determine how this information may be used to provide further inputs into the decision-making process of the team coordination agent. For example, the content of voice communications can be analyzed to detect and refine representations of operator beliefs, desires, and intentions. Likewise, vocal content and tone could also be used to measure operator state, such as stress (Van Puyvelde et al., 2018).

Like Phase 2, by integrating teams in Phase 3, ML approaches can be leveraged to train more realistic Phase 1/2 synthetic agents to behave in patterns drawn from the behavior of actual human teams. This can include refining interactions with synthetic squad members during Phase 2 evaluations. While all phases require some objective metric to assess the performance of the AI capability, team-level performance metrics may be particularly critical for demonstrating gains obtained with the AI capability. Early Phase 3 evaluations using heuristic approaches can be used to refine the content of objective team performance metrics. Teamwork performance metrics can be difficult to define and evaluate due to the complex cognitive and social components that lead to effective teamwork. Through an iterative evaluation of the system with stakeholders and test participants, task relevant teamwork performance metrics could be refined in conjunction with the design of the entire system. These performance metrics could also be used to repurpose the system for mission training, such as being implemented as an external assessment engine in the Generalized Intelligent Framework for Tutoring (Sottolare et al., 2012; Vatrak et al., 2022).

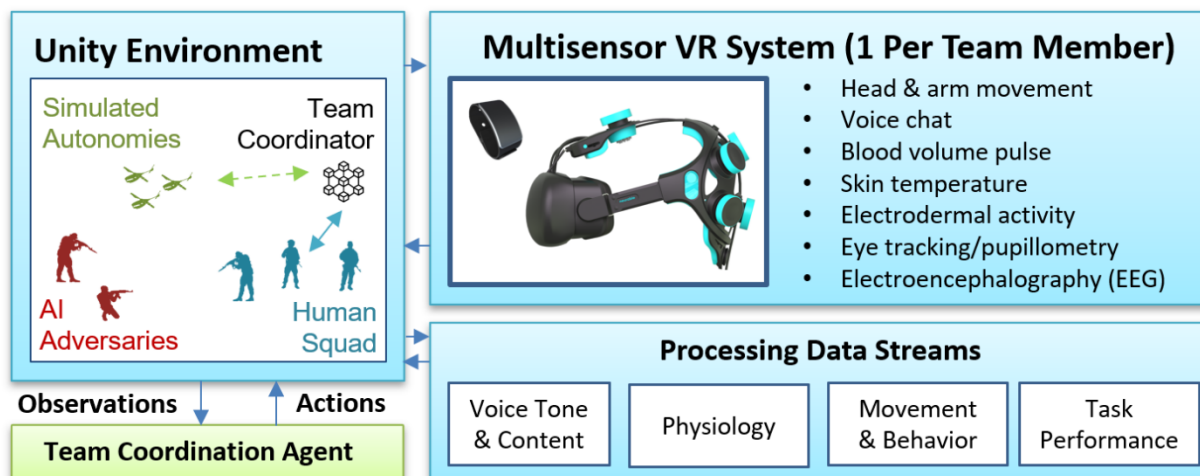


Figure 6. Phase 3 overview. A team of humans evaluate a team coordination agent coordinating a multi-human multi-agent team in a simulated operational scenario. Human participants' behavior and physiology are monitored using a VR system with integrated sensors. Realtime processing of physiological and behavioral data can be fed into the team coordination agent's observations. This data can further refine the agent's recommended actions to meet the needs of both individuals and the team, as well as objectively evaluate the impact of the system on operators. Image of VR headset with integrated EEG provided by Wearable Sensing (San Diego, CA).

DISCUSSION

While an AiTR-AR system can improve individual task performance, further performance gains may be realized by coordinating target allocation at the squad-level. When reviewing the performance of a rifle-mounted AiTR-AR system, evaluation participants and stakeholders suggested ways a team coordination agent could improve squad coordination. Coordinating target allocation through an IDA could improve squad lethality while reducing cognitive burden by focusing attention on assigned targets. IDAs that track user capabilities, behaviors, and states (e.g., intent to engage a target, taking enemy fire) could further enhance teamwork performance. Human aware planning applied to multi-human multi-agent teams provides a general framework for designing autonomies that can infer both operator and adversarial intent, develop a joint plan, and provide human operators with explainable planning decisions.

We defined a three-stage approach to develop and evaluate a team coordination AI for multi-human multi-agent teams. While we are using this framework to design a team coordination AI for the dismounted soldier through a rifle-mounted AR display, this approach is applicable to any teaming capability that can be simulated and tested in VR. By leveraging VR, the AI UI/UX can be evaluated with human operators early in development process. Heuristics that leverage absolute knowledge of the virtual environment can be used to evaluate potential performance gains before any AI model training is necessary. Human- and team-in-the-loop evaluations can be used to both refine system user needs as well as synthetic human model behaviors utilized for training.

We utilized Unity as a general-purposes development environment as it enables complex agent-based simulations, including a comprehensive and extensible toolkit for reinforcement learning and ML model integration. Unity also allows for platform-independent VR, physiological sensor integration, and multi-client networking. Unity-based ML agent training has previously been utilized in the development of adaptive synthetic characters in dismounted soldier simulations for training (Liu et al., 2021; Ustun et al., 2020). In contrast, our primary objective is to utilize reinforcement learning to improve team coordination, although ML approaches may also be utilized in our approach to improve synthetic human behaviors. Our approach utilizes early human- and team-in-the-loop evaluations to improve synthetic agent behaviors to provide a more representative proxy during team coordination agent training. This can include straightforward measures (e.g., reaction times), as well as more subtle behaviors that may be enhanced from relatively few training samples such as verbal (Reed et al., 2021) and nonverbal (Punzi et al., 2022) communication. Our approach is complementary to emerging human-machine teaming design frameworks based around user-centric design and validation (Dominguez et al., 2020; McDermott et al., 2018; Scielzo et al., 2021).

CONCLUSION

We describe an iterative three-phase framework to aid development of teamwork-promoting autonomous capabilities:

- Phase 1: Synthetic agents approximate human behaviors to train AI models for team coordination.
- Phase 2: Individual human operators evaluate the UI/UX in the simulated operational environment.
- Phase 3: Humans teams evaluate how the capability may complement and enhance teamwork performance.

Qualitative and quantitative feedback from each phase can be used to improve simulation fidelity, AI performance, and the UI/UX of the system. Key benefits of this approach are to enable:

- Efficient use of development resources via a continuous workflow from ML to human-in-the-loop evaluation.
- Iterative, human-centered design to enhance autonomy performance and teammate acceptance.
- A VR system that can be used for the training and evaluation of future teams to interact with the autonomy.

ACKNOWLEDGEMENTS

This research was supported by Riverside Research Independent Research and Development funds.

REFERENCES

- Bansal, G., Wu, T., Zhou, J., Fok, R., Nushi, B., Kamar, E., Ribeiro, M. T., & Weld, D. (2021). Does the Whole Exceed its Parts? The Effect of AI Explanations on Complementary Team Performance. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 1–16. <https://doi.org/10.1145/3411764.3445717>
- Bobb, R. L., Coady, J. A., Barnard, V. O., Mueller, M. A., Casebeer, W. D., & Salisbury, J. P. (2022). Virtual reality framework for design and evaluation of multispectral computer vision algorithms and augmented reality

- interfaces for enhancing situational awareness in degraded visual environments. *Virtual, Augmented, and Mixed Reality (XR) Technology for Multi-Domain Operations III*, 13.
- Brockman, G., Cheung, V., Pettersson, L., Schneider, J., Schulman, J., Tang, J., & Zaremba, W. (2016). OpenAI Gym. *ArXiv:1606.01540 [Cs]*. <http://arxiv.org/abs/1606.01540>
- Cao, Z., Wong, K., Bai, Q., & Lin, C.-T. (2020). Hierarchical and Non-Hierarchical Multi-Agent Interactions Based on Unity Reinforcement Learning. *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS*, 3.
- Chakraborti, T., Sreedharan, S., & Kambhampati, S. (2018). Human-Aware Planning Revisited: A Tale of Three Models. *IJCAI-ECAI XAI/ICAPS XAIP Workshops*, 10.
- DeCostanza, A. H., Marathe, A. R., Bohannon, A., Evans, A. W., Palazzolo, E. T., Metcalfe, J. S., & McDowell, K. (2018). *Enhancing Human-Agent Teaming with Individualized, Adaptive Technologies: A Discussion of Critical Scientific Questions*. US Army Research Laboratory Aberdeen Proving Ground United States. <https://apps.dtic.mil/sti/citations/AD1051552>
- Dominguez, D. C. O., Bonaceto, C. A., Johnson, C., Bryant, M., Rahman, M., & Darling, E. (2020). *Cognitive Engineering Toolkit*. <https://www.mitre.org/publications/technical-papers/cognitive-engineering-toolkit>
- Frame, M. E., Boydston, A. S., & Lopez, J. S. (2020). Development of an Autonomous Manager for Dynamic Human-Machine Task Allocation in Operational Surveillance. *2020 IEEE International Conference on Human-Machine Systems (ICHMS)*, 1–4. <https://doi.org/10.1109/ICHMS49158.2020.9209414>
- Geuss, M. N., Larkin, G., Swoboda, J., Yu, A., Bakdash, J., White, T., Berthiaume, M., & Lance, B. (2019). Intelligent squad weapon: Challenges to displaying and interacting with artificial intelligence in small arms weapon systems. *Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications, 11006*, 296–307. <https://doi.org/10.1117/12.2518405>
- Grimm, D. A., Demir, M., Gorman, J. C., & Cooke, N. J. (2018). Team Situation Awareness in Human-Autonomy Teaming: A Systems Level Approach. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 62(1), 149–149. <https://doi.org/10.1177/1541931218621034>
- Hoffman, M., Shahriari, B., Aslanides, J., Barth-Marion, G., Behbahani, F., Norman, T., Abdolmaleki, A., Cassirer, A., Yang, F., Baumli, K., Henderson, S., Novikov, A., Colmenarejo, S. G., Cabi, S., Gulcehre, C., Paine, T. L., Cowie, A., Wang, Z., Piot, B., & de Freitas, N. (2020). Acme: A Research Framework for Distributed Reinforcement Learning. *ArXiv:2006.00979 [Cs]*. <http://arxiv.org/abs/2006.00979>
- Juliani, A., Berges, V.-P., Teng, E., Cohen, A., Harper, J., Elion, C., Goy, C., Gao, Y., Henry, H., Mattar, M., & Lange, D. (2020). Unity: A General Platform for Intelligent Agents. *ArXiv:1809.02627 [Cs, Stat]*. <http://arxiv.org/abs/1809.02627>
- Kozlowski, S. W. J., & Klein, K. J. (2000). A multilevel approach to theory and research in organizations: Contextual, temporal, and emergent processes. In *Multilevel theory, research, and methods in organizations: Foundations, extensions, and new directions* (pp. 3–90). Jossey-Bass.
- Lance, B. J., Larkin, G. B., Touryan, J. O., Rexwinkle, J. T., Gutstein, S. M., Gordon, S. M., Toulson, O., Choi, J., Mahdi, A., Hung, C. P., & Lawhern, V. J. (2020). Minimizing data requirements for soldier-interactive AI/ML applications through opportunistic sensing. *Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications II, 11413*, 19–27. <https://doi.org/10.1117/12.2564515>
- Larkin, G. B., Geuss, M., Yu, A., Rexwinkle, J., Callahan-Flintoft, C., Bakdash, J. Z., Swoboda, J., Lieberman, G., Hung, C. P., & Moore, S. (2020). Augmented target recognition display recommendations. *DSIAC Journal*, 7(1), 28–34.
- Larkin, G., Geuss, M., Yu, A., Callahan-Flintoft, C., Rexwinkle, J., Hung, C., & Lance, B. (2020). *Human Performance with Complex Technology: How Visual Cognition Is Critical to Enhanced Performance with Aided Target Recognition (AiTR)* (pp. 134–140). https://doi.org/10.1007/978-3-030-51041-1_19
- Larson, L., & DeChurch, L. A. (2020). Leading teams in the digital age: Four perspectives on technology and what they mean for leading teams. *The Leadership Quarterly*, 31(1), 101377. <https://doi.org/10.1016/j.leaqua.2019.101377>
- Lee, M., Tweedell, A., Dennison, M. S., Sabbagh, P., Conroy, J., Trout, T., Freeman, J., & Lance, B. (2022). Modeling and analysis of motion data from dynamic soldier state estimation to enable situational understanding. *Virtual, Augmented, and Mixed Reality (XR) Technology for Multi-Domain Operations III*, 13. <https://doi.org/10.1117/12.2618634>
- Liu, L., Gurney, N., McCullough, K., & Ustun, V. (2021). Graph Neural Network Based Behavior Prediction to Support Multi-Agent Reinforcement Learning in Military Training Simulations. *2021 Winter Simulation Conference (WSC)*, 1–12. <https://doi.org/10.1109/WSC52266.2021.9715433>

- Madni, A. M., & Madni, C. C. (2018). Architectural Framework for Exploring Adaptive Human-Machine Teaming Options in Simulated Dynamic Environments. *Systems*, 6(4), 44. <https://doi.org/10.3390/systems6040044>
- Mait, J., Poree, D., Prater, J., Reynolds, P., Stepp, D., West, B. J., Kott, A., Swami, A., DeCostanza, A. H., Franaszczuk, P. J., McDowell, K., Piekarski, B., Sadler, B. M., Carter, R., & Zabinski, J. (2017). *2015 Army Science Planning and Strategy Meeting Series: Outcomes and Conclusions*. 90.
- Marks, M. A., Mathieu, J. E., & Zaccaro, S. J. (2001). A Temporally Based Framework and Taxonomy of Team Processes. *Academy of Management Review*, 26(3), 356–376. <https://doi.org/10.5465/amr.2001.4845785>
- McDermott, P., Dominguez, C., Kasdaglis, N., Ryan, M., Trhan, I., & Nelson, A. (2018). *Human-machine teaming systems engineering guide*. MITRE CORP BEDFORD MA BEDFORD United States.
- O'Neill, T., McNeese, N., Barron, A., & Schelble, B. (2020). Human–Autonomy Teaming: A Review and Analysis of the Empirical Literature. *Human Factors*. <https://doi.org/10.1177/0018720820960865>
- Punzi, M., Ladeveze, N., Nguyen, H., & Ravenet, B. (2022). ImCasting: Nonverbal Behaviour Reinforcement Learning of Virtual Humans through Adaptive Immersive Game. *27th International Conference on Intelligent User Interfaces*, 62–65.
- Rabinowitz, N. C., Perbet, F., Song, H. F., Zhang, C., Eslami, S. M. A., & Botvinick, M. (2018). Machine Theory of Mind. *ArXiv:1802.07740 [Cs]*. <http://arxiv.org/abs/1802.07740>
- Raffin, A., Hill, A., Gleave, A., Kanervisto, A., Ernestus, M., & Dormann, N. (2021). Stable-Baselines3: Reliable Reinforcement Learning Implementations. *Journal of Machine Learning Research*, 8.
- Reed, L., Li, C., Ramirez, A., Wu, L., & Walker, M. (2021). *Jurassic is (almost) All You Need: Few-Shot Meaning-to-Text Generation for Open-Domain Dialogue* (arXiv:2110.08094). arXiv. <http://arxiv.org/abs/2110.08094>
- Roth, E. M., Sushereba, C., Militello, L. G., Diiulio, J., & Ernst, K. (2019). Function Allocation Considerations in the Era of Human Autonomy Teaming. *Journal of Cognitive Engineering and Decision Making*, 13(4), 199–220. <https://doi.org/10.1177/1555343419878038>
- Salas, E., Stagl, K. C., Burke, C. S., & Goodwin, G. F. (2007). Fostering team effectiveness in organizations: Toward an integrative theoretical framework. In *Modeling complex systems* (pp. 185–243). University of Nebraska Press.
- Schaefer, K. E., Perelman, B., Rexwinkle, J., Canady, J., Neubauer, C., Waytowich, N., Larkin, G., Cox, K., Geuss, M., Gremillion, G., Metcalfe, J. S., DeCostanza, A., & Marathe, A. (2021). Human-Autonomy Teaming for the Tactical Edge: The Importance of Humans in Artificial Intelligence Research and Development. In *Systems Engineering and Artificial Intelligence* (pp. 115–148). Springer International Publishing. https://doi.org/10.1007/978-3-030-77283-3_7
- Scielzo, S., Kocak, D. M., & Technologies, H. (2021). A Multi-Domain Robotic Teammate Framework: Next Generation Human-Machine Interface Guidelines to Support Trust and Mission Outcomes. *IITSEC*, 13.
- Scott, N. V., Martinez, L. M., & Salisbury, J. P. (2021, November). *Application of Tiny YOLOv2 Car Detection for Traffic Flow Quantification during the COVID Health Pandemic: Preliminary Results*. 2nd International Conference on Civil Engineering Fundamentals and Applications (ICCEFA'21). <https://doi.org/10.11159/iccefa21.126>
- Seeber, I., Bittner, E., Briggs, R. O., de Vreede, T., de Vreede, G.-J., Elkins, A., Maier, R., Merz, A. B., Oeste-Reiß, S., Randrup, N., Schwabe, G., & Söllner, M. (2020). Machines as teammates: A research agenda on AI in team collaboration. *Information & Management*, 57(2), 103174. <https://doi.org/10.1016/j.im.2019.103174>
- Sengupta, S., Chakraborti, T., & Kambhampati, S. (2018). Ma-radar—a mixed-reality interface for collaborative decision making. *ICAPS UISP*.
- Sottolare, R. A., Brawner, K. W., Goldberg, B. S., & Holden, H. K. (2012). The generalized intelligent framework for tutoring (GIFT). *Orlando, FL: US Army Research Laboratory–Human Research & Engineering Directorate (ARL-HRED)*.
- Ustun, V., Kumar, R., Reilly, A., Sajjadi, S., & Miller, A. (2020). Adaptive Synthetic Characters for Military Training. *IITSEC*, 11.
- Van Puyvelde, M., Neyt, X., McGlone, F., & Pattyn, N. (2018). Voice Stress Analysis: A New Framework for Voice and Effort in Human Performance. *Frontiers in Psychology*, 9. <https://www.frontiersin.org/article/10.3389/fpsyg.2018.01994>
- Vatral, C., Mohammed, N., Biswas, G., & Goldberg, B. (2022). Moving Beyond Training Doctrine to Explainable Evaluations of Teamwork using Distributed Cognition. *Generalized Intelligent Framework for Tutoring (GIFT) Users Symposium (GIFTSym10)*, 127.
- Williams, J., Fiore, S. M., & Jentsch, F. (2022). Supporting Artificial Social Intelligence With Theory of Mind. *Frontiers in Artificial Intelligence*, 5, 750763. <https://doi.org/10.3389/frai.2022.750763>