



ARL-TR-9435 • APR 2022



Distributed and Collaborative Intelligent Systems and Technology (DCIST) Collaborative Research Alliance DCIST CRA W911NF-17-2- 0181 2018–2020 (Summary Technical Report, Oct 2018–Sep 2020)

**by Brett Piekarski, Brian M Sadler, Vijay Kumar, and
Alejandro Ribeiro**

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.



**Distributed and Collaborative Intelligent Systems
and Technology (DCIST) Collaborative Research
Alliance DCIST CRA W911NF-17-2-0181 2018–2020
(Summary Technical Report, Oct 2018–Sep 2020)**

Brett Piekarski and Brian M Sadler
DEVCOM Army Research Laboratory

Vijay Kumar and Alejandro Ribeiro
University of Pennsylvania

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) April 2022		2. REPORT TYPE Summary Technical Report		3. DATES COVERED (From - To) 1 October 2018–30 September 2020	
4. TITLE AND SUBTITLE Distributed and Collaborative Intelligent Systems and Technology (DCIST) Collaborative Research Alliance DCIST CRA W911NF-17-2-0181 2018–2020 (Summary Technical Report, Oct 2018–Sep 2020)				5a. CONTRACT NUMBER DCIST CRA W911NF-17-2-0181	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Brett Piekarski, Brian M Sadler, Vijay Kumar, and Alejandro Ribeiro				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) DEVCOM Army Research Laboratory ATTN: FCDD-RLC Aberdeen Proving Ground, MD 21005				8. PERFORMING ORGANIZATION REPORT NUMBER ARL-TR-9435	
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 Approved for public release: distribution unlimited.					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT The Distributed and Collaborative Intelligent Systems and Technology (DCIST) Collaborative Research Alliance (CRA) is exploring key underlying technologies and methods needed to enable robotics and autonomous systems operations across all Army-relevant environments; provide better situational awareness; increase warfighter capabilities and standoff; increase coverage and create dilemmas for the adversary; provide force multiplication; enable faster decision-making; and extend maneuverability in ways yet to be imagined. The DCIST CRA is advancing the state of the art in distributed intelligence, heterogeneous group control, and adaptive and resilient behaviors.					
15. SUBJECT TERMS autonomy, autonomous systems, robotics, collaborative robotics, distributed perception and control, Military Information Sciences, STR					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 123	19a. NAME OF RESPONSIBLE PERSON Brett Piekarski
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified			19b. TELEPHONE NUMBER (Include area code) 240-618-0120

Contents

List of Figures	iv
Acknowledgments	viii
Executive Summary	ix
1. Introduction	1
2. DCIST Research Plan	3
3. Methods and Findings: Progress to Date and Future Plans	5
3.1 Thrust 1: Distributed Intelligence	5
3.1.1 Capabilities Description and Major Accomplishments	7
3.2 Thrust 2: Heterogeneous Group Control	26
3.2.1 Capabilities Description and Major Accomplishments	27
3.3 Thrust 3: Adaptive and Resilient Behaviors	38
3.3.1 Capabilities Description and Major Accomplishments	40
3.4 Thrust 4: Cross Disciplinary Experiments (CDE)	58
3.4.1 Capabilities Description	58
4. Key Program Metrics and Impact (18Q3 to 20Q3)	65
4.1 Education: Number of Students Supported	65
4.2 Number of Publications	65
4.3 Notable Papers, Awards, and Recognitions	65
5. Conclusion	66
6. References	68
Appendix. DCIST Program Notable Publications (18Q3–21Q3)	73
List of Symbols, Abbreviations, and Acronyms	105
Distribution List	107

List of Figures

Fig. 1	Dense metric-semantic 3-D mesh model of an urban scene generated by Kimera-Multi with three robots	8
Fig. 2	An aerial vehicle in a collaborative forest mapping mission	9
Fig. 3	Fused 3-D occupancy grid of environment and vehicle trajectories at the end of the mission. Vehicle paths are shown in blue and yellow and each coordinate frame represents the origin of a submap	10
Fig. 4	Learning robotic navigation: ViNG builds and plans over a learned topological graph consisting of previously seen egocentric images, and uses a learned controller to execute the path to a visually indicated goal. Unlike prior work, our method uses purely offline experience and does not require a simulator or online data collection. Note that the graph constructed by our algorithm is not geometric and nodes are not associated with coordinates in the world, but only with image observations—the top-down satellite image is provided only for visualization and is not available to our method	12
Fig. 5	Cooperative multi-agent tasks: Overview of grid-world environments used in our experiments. Cooperative and self-interested agents are visualized as blue and red squares, respectively. Black cells correspond to obstacles. In the coverage environments, different colors indicate the coverage achieved by individual agents. In the path planning environment, labeled goal locations are indicated by circles.	13
Fig. 6	Optimization landscape of single-task/multi-task RL. Unlike single-task learning problems, multi-task RL problems can have isolated local minima, meaning that there are no a priori guarantees that descent algorithms will find globally optimal solutions	15
Fig. 7	MSF for (simulated) aerial drones trained to navigate in different indoor environments. In the table on the right SA-i is a drone trained only in environment i ; the policy learned performs well in environment i but poorly in the other three environments. When a single policy is learned for all four environments using distributed optimization (row DCPG), the results are superior across all cases	15
Fig. 8	(a) $-G(V,E)$ and its evolution as new nodes and edges are added—the goal is to update the solution to the associated optimization program efficiently. (b) We reduce the general graph (left) to the chain graph (right); starting from the new node f_8 , expanding one edge outward each time, we collect all unassigned neighbors into a single node	17
Fig. 9	Planning results for various abstraction levels	19
Fig. 10	Robot trajectories that satisfy the LTL task. The robots are initially located at the top left corner. The green and red squares represent the initial and final locations of the targets while the black paths correspond to the target trajectories	21

Fig. 11	Perimeter defense: A team of agents defends a perimeter	23
Fig. 12	Flocking tasks to highlight the ability of our approach to handle dynamic communication networks	24
Fig. 13	Encoding the task as a graph: the known map locations and team members are graph nodes, and allowed moves are graph edges.....	24
Fig. 14	Demonstrating the scalability of our methods for a large number of robots.....	25
Fig. 15	Redundant assignment answers the question, which task should receive extra resources? Given an existing non-redundant assignment and respective cost distributions (top right), a single redundant robot, Robot 2, can be assigned to either task. Utilitarian approaches assign Robot 2 to Person 0 because of the higher improvement in cost (Graph A), whereas fair approaches assign Robot 2 to Person 1 because of their higher need (Graph B).	30
Fig. 16	By learning, in real-time, models of environmental disturbances and individual agents' suitability to particular tasks, an adaptive and robust framework is obtained for task allocation and execution for heterogeneous teams.	31
Fig. 17	Overview of the proposed two-level approach to learning heterogeneous multi-robot coordination	32
Fig. 18	A limited-sensing model for target guarding	33
Fig. 19	The red triangles represent the attackers, and the blue circle denotes the defender. At least three defenders are required to neutralize attack.	34
Fig. 20	Principles of design for optimal communication in multi-agent teams under constraints of time and limited sensing.....	36
Fig. 21	The Network Inference pretest (left) is a simple 5-min cognitive test. Human operator performance on this pretest is highly correlated (middle) with operator performance on complex multi-robot teaming tasks, such as teleoperating multiple ground and aerial vehicles while maintaining an ad-hoc communication network (right). The ability to quickly determine which human operators are better suited for robot control can help inform task allocation in complex HRT scenarios...	37
Fig. 22	Decomposition of global active information acquisition tasks into local tasks.....	42
Fig. 23	Codes for multi-agent reinforcement learning	43
Fig. 24	Large-scale experimental setup to test our custom infrastructure	45
Fig. 25	Designing coordination strategies to allow robots to periodically encounter at predetermined location	46
Fig. 26	Developing a locally optimal solution in three steps	47
Fig. 27	A team of robots (green unmanned aerial vehicles [UAVs]) is tasked to track multiple targets (red UAVs). The adversary can attack the	

	sensing/observations (black arrows) of the robots and blank communications (blue dotted lines) between them. The objective is to minimize the uncertainty (light blue ellipses) of the robots' positions despite the sensing and communication attacks.....	48
Fig. 28	Hostile target tracking using a proximity-based model	49
Fig. 29	Experiments with a resilient networked testbed: A team of five quadrotors executing coordinated motion are monitored by ground robots.....	50
Fig. 30	Example of a heterogeneous team tracking the movement of a herd of animals	52
Fig. 31	Learning to swarm using K-NODEs.....	53
Fig. 32	Example of different paths with different detection likelihoods taking into consideration terrain geometry	54
Fig. 33	Quantitative comparisons of two variants of model-based meta-RL (McAllister 2021) developed as part of this project (gradient-based adaptive learner [GrBAL] and recurrence-based adaptive learner [ReBAL] + model predictive path integral control [MPPI]) on simulated (left) and real-world (middle) benchmark tasks that require rapid adaptation to test-time changes, including mechanical damage to a robot and changing terrains. Right image shows the real-world legged platform used for experiments. The y-axis shows (normalized) total reward over the course of an episode (which includes adaptation time), after meta-trained for a fixed number of samples (chosen to be realistically low, equivalent to a few hours in simulated tasks and 30 min for the real-world tasks). Our methods (GrBAL and ReBAL) attain results that are significantly better than prior approaches, often by 2 $\times$ or more.....	56
Fig. 34	Quantitative comparison of our method, PEARL, with state-of-the-art meta-RL algorithms circa 2019 at the time our method was developed. The x-axis is a log scale showing the number of samples needed by the algorithm, and the y-axis is the performance (total reward) at that point during meta-training on unseen test tasks. Across most benchmarks, PEARL learns about two orders of magnitude faster, and often attains significantly better final results.....	57
Fig. 35	A team is moving through a contested urban area, with multiple routes through the environment and where dynamic agents are to be avoided	59
Fig. 36	Two example trajectories of a quadrotor choosing minimum visibility trajectories that it has learned to predict from the semantics in the environment	62
Fig. 37	The test facility at Medfield, Massachusetts, outside Boston. The blue trajectory is the initial trajectory of a Clearpath Jackal, and the green trajectory is the initial trajectory of a quadrotor using the stealthy semantic navigation strategy. When objects of interests are detected by the Jackal, the quadrotor is tasked with deviations to inspect and refine	

	the pose of the objects, denoted by the dashed green lines. This scenario exercises the state estimation, object localization, coordinate frame alignment, and stealth navigation.	62
Fig. 38	(Left) Patrolling operations of 15 Warthog robots in the DCIST Unity simulation. (Right) Perimeter defense operations on two urban perimeters using six Warthog defender agents (blue) and six aerial intruders (red).....	63
Fig. 39	(Left) shows a configuration of aerial robots that are dedicated to ensuring communication among the defense team. These agents leverage a connectivity maintenance algorithm to constantly move in order to maintain connectivity among the task agents. (Right) shows the view of a building from a quadroter.	64

Acknowledgments

The following DCIST CRA Thrust Area and Capability Task Leads contributed the “Capabilities Description and Major Accomplishments” write-ups for each thrust and capability captured in Section 3 of this report:

- Ethan Stump, Thrust 1 Co-Lead, DEVCOM ARL
- Nick Roy, Thrust 1 Co-Lead, Massachusetts Institute of Technology
- Christopher Kroninger, Thrust 2 Co-Lead, DEVCOM ARL
- Magnus Egerstedt, Thrust 2 Co-Lead, Georgia Institute of Technology
- Michael Dorothy, Thrust 2 Deputy Lead, DEVCOM ARL
- Jonathan Fink, Thrust 3 Co-Lead, DEVCOM ARL
- George Pappas, Thrust 3 Co-Lead, University of Pennsylvania
- Jonathan How, Capability Lead, Massachusetts Institute of Technology
- Amanda Prorok, Capability Lead, University of Cambridge
- Justin Romberg, Capability Lead, Georgia Institute of Technology
- Gaurav Sukhatme, Capability Lead, University of Southern California
- Sonia Chernova, Capability Lead, Georgia Institute of Technology
- M Ani Hsieh, Capability Lead, University of Pennsylvania
- Sergey Levine, Capability Lead, University of California, Berkeley
- Dilip Patel, Koniag Information Security Services, LLC

Executive Summary

As the 2020 Nagorno-Karabakh war demonstrated, the role of drones is becoming increasingly critical and will be one of the key decisive factors in future operations. Extrapolating the number and type of autonomous platforms (both ground and aerial Robotics and Autonomous Systems [RAS]) and considering near-peer adversaries motivate the Distributed and Collaborative Intelligent Systems and Technology (DCIST) Collaborative Research Alliance (CRA). This innovative program brings together leading RAS academic and US Army Combat Capabilities Development Command Army Research Laboratory researchers with the objective of significantly increasing our capability to scale RAS technologies, to prevent adversarial advantage, and to counter and neutralize an adverse force. The DCIST CRA spirals out technologies, shows experimental proofs for others to build upon, and trains the next generation of Army researchers that will truly enable RAS to become integrated into the Army of the future. This technical report summarizes the technological challenges and research progress through September 2020.

Through basic research, DCIST is exploring key underlying technologies and methods needed to enable RAS operations across all Army-relevant environments; provide better situational awareness; increase warfighter capabilities and standoff; increase coverage and create dilemmas for the adversary; provide force multiplication; enable faster decision-making; and extend maneuverability in ways yet to be imagined. The program is organized into the following three capability-based research thrusts plus a set of cross-disciplinary experiments:

Distributed Intelligence: This thrust focuses on multi-agent intelligent system perception, communication, and planning. These systems must plan and work together efficiently, learn collaboratively, and adapt to wireless networking. Specific research goals include the following:

- Develop knowledge representations that enable adaptive perception-action-communication loops for distributed intelligence.
- Establish a theory of hierarchical and abstract representations that enable distributed inference and decision-making.
- Enable efficient learning across scalable teams that is modular, composable, and data-efficient.

Heterogeneous Group Control: This thrust focuses on control of large autonomous teams with varying levels of heterogeneity and degrees of autonomy. Advancing the state of the art in this thrust requires developing fundamental understanding and formalism necessary to incorporate a variety of

robots within operational and tactical teams. The work is capability driven, to enable hierarchical and distributed control for adversarial operations, scalable task assignment for heterogeneous multi-unit teams, tactical engagement in complex environments, and Soldier interaction. Research goals include the following:

- Control and communication for scalable teams
- Tactical heterogeneous team behaviors
- Multi-agent coordination with Soldiers

Adaptive and Resilient Behaviors: This thrust focuses on robustness and adaptation of heterogeneous teams to achieve resilience to failures, loss, and degraded communications in the face of dynamic intelligent adversaries and changing environmental conditions. Some specific areas of research include resilient situational awareness, wireless communication networks for distributed collaborative systems, exploiting heterogeneity for resilience, multi-agent behaviors in the presence of adversaries, and rapid adaptation to large disturbances. The research in this thrust addresses two key questions:

- How do we build algorithms and deploy scalable systems that can maintain effectiveness despite operation in heavily contested situations?
- How do large distributed intelligent systems react to rapid change due to an adversary without sacrificing stability or performance?

Cross-Disciplinary Research Experiments (CDEs): The CDEs are designed to bring multiple principal investigators together in an experimental setting, explore and discover interdependencies across research areas, and experimentally demonstrate new capabilities. The CDEs include physical multi-robot experiments and a 3-D realistic simulation environment. The two CDEs are structured around

- Heterogeneous Multi-Agent Situational Awareness
- Dynamic Teaming Operations in Contested Environments

The DCIST team is exploring and developing novel algorithms in artificial intelligence (AI)/machine learning (ML), autonomy, and robotics. These take multi-sensory input from multiple heterogeneous robots (ground and air) to perceive the environment and aid in system- and subsystem-level autonomous team operation in contested environments. A key aspect of the algorithm effort is to leverage advances in AI to enable robust collaborative autonomous robotic operations and behaviors. Some examples are learning communications for multi-

robot planning, learning optimal resource allocation in multi-agent systems, semantic scene understanding and inference for building a rich perception of the world, learning a sparse representation of the environment that is sufficiently lightweight for long distance path planning and navigation, learning decentralized Perception Action Communication loops for robot teams, and integrating learning and planning for resilient autonomous behaviors in complex and adversarial environments.

The DCIST CRA also addresses technology enablers, such as a robust and resilient wireless networking, and data architectures. This ensures that Army-relevant communication constraints are always considered and accounted for, including denied, intermittent, or limited communications. DCIST is exploring and developing novel robust/resilient algorithms to optimize communications of heterogeneous team members, use mobile robotics to adapt and maintain communications in distributed operations, address information sharing in intermittent communications environments, and learning what and with whom to communicate in order to adapt and optimize resources and behaviors. This research is essential for the DCIST program's cross-disciplinary experiments where multiple heterogeneous agents have to rely on a network to reliably exchange data. Likewise, DCIST is focused on extending RAS perception, planning, and tactical behaviors to scalable heterogeneous teams in Army-relevant operational environments and while under realistic infrastructure constraints (e.g., no GPS, stale maps, limited/degraded communications) and mission, environmental, and adversarial complexity.

Selected significant accomplishments by thrust area through the first three years of DCIST are as follows:

Distributed Intelligence

- A framework for multi-modal and hierarchical knowledge representation that includes both metric and semantic information. A specific example is “Kimera-Multi,” which is a distributed multi-robot mapping system that builds a metric-semantic model of an environment in real time and with limited communication.
- Methods that learn in a way that generalizes to multiple tasks and objectives and can learn skills and capabilities in the context and constraints of a team to solve heterogeneous multi-agent coordination problems.
- A framework for vision-based robotic navigation that can readily generalize to diverse real-world environments, operate in GPS-denied environments, and enable simple and accessible goal specification by humans or upstream

planning algorithms through desired locations, images of goal landmarks, and latent (learned) goal representations.

- A mathematical and algorithmic framework for distributed problems that respect communication constraints where autonomous agents are expected to work together to solve inference, learning, and control problems.
- Planning models that are hierarchical and composable, and that represent both prior and acquired knowledge to deal with complexity and allow more efficient planning.
- Advancement of graph neural networks as a fundamental enabling technology for designing agent controllers that work for large-scale collaborative systems.

Heterogeneous Group Control

- Theory and algorithms aimed at strategic deployment of agents in large environments with dynamically changing scenarios involving models of adversarial agents and imperfect/delayed communication when confronted with the realities of fast-moving swarm-versus-swarm engagements and the presence of imperfect information exchange and nonrational players.
- Methods to produce sufficient or “good enough” solutions in a computationally feasible, distributed, and adaptive manner to rapidly adapt and retask to ensure mission success (i.e., not waiting for “optimal” solutions that may not be achievable for the given task and context).
- Methods and tools for the analysis and development of cooperative team strategies in the presence of adversarial agents.
- Improved modeling techniques for how humans perceive and communicate time-evolving information, and frameworks for modeling human capabilities with the goal of improving human–robot coordination in the context of teams operating in complex environments.

Adaptive and Resilient Behaviors

- Active sensing techniques that go beyond optimization for computational efficiency and information gain that dynamically adjust their perception and motion to achieve resilient situational awareness in the presence of sensor failures, jammed communications, detection risk, and/or compromised agents.
- Methods that dynamically adjust to achieve resilient communication and mobility at scale for teams of heterogeneous agents that ensures mission

progress and preserves core capabilities of the team in the face of failures, disruption, and loss.

- Multi-agent reinforcement learning algorithms that can empower agents with resilience in dynamic environments or in the presence of adversaries who may rely on deception or other intelligent strategies at mission time.
- Tools to track and engage an adversarial team while managing uncertainty in strategy, execution, and environment for multi-robot surveillance and perimeter security scenarios.
- Adaptive ML algorithms that enable teams of robots to adapt on-the-fly to unforeseen large and rapid changes in environmental conditions, mission parameters, and robot state by means of online learning and meta-learning algorithms.

Details on these and other advancements can be reviewed in Section 3, “Methods and Findings: Progress to Date and Future Plans,” and a full bibliography for the period through FY20 is included in the Appendix.

1. Introduction

Future Multi-Domain Operations (MDO) will require a Force that can conduct cross-domain strategic maneuvers to penetrate and operate in complex and contested areas. The operational concepts within MDO rely heavily on advancements in Robotic and Autonomous Systems (RAS). Hence, RAS are being explored to operate across all environments to provide better situational awareness, increase warfighter capabilities and standoff, increase coverage and dilemmas to the adversary, allow for Force multiplication, and provide the ability to greatly extend maneuverability in ways yet-to-be-imagined to support reconnaissance, breach, attack, protect, and sustainment operations. To realize these scalable concepts within real teams of heterogeneous RAS and in Army relevant environments, significant research gaps must be overcome. For example, these scalable RAS teams will need the capability to do the following:

- Assess dynamic scenes and create and share information in a scalable way across heterogeneous teams within the communication resources available at that point in time.
- Coordinate across echelons, teams, sub-teams, and individual systems to operate tactically and create windows of opportunities and overmatch situations.
- Collaboratively perceive and learn the context of the operational environment to understand, model, predict, and adapt to adversary force maneuvers.
- Adapt and be resilient to large disturbances in understanding, changes in the environment or available infrastructure, and adversarial operations to maintain mission success.

Several axes of complexity exist that limit these operational capabilities of RAS in relevant Army environments and operations and therefore must also be addressed. These include the following:

- **Mission complexity**, including scalability in physical area and number of platforms, risk, duration, numbers and types of tasks, operational tempo, and many other factors.
- **Environmental complexity**, including all forms of man-made, man-destroyed, and natural terrain.
- **Adversarial complexity**, including counter-autonomy in its many potential forms.

Environments, such as complex off-road terrain and dense urban, pose severe challenges to mobility, perception, networking, and sensing for both ground and small aerial platforms operating within the environment. Limited prior access to the operational environment may limit the application of today's big-data learning approaches. The complexity of human behavior makes discerning noteworthy behaviors or formulating appropriate responses challenging. And there may be little or no available infrastructure, such as power, GPS, or communication networks. For the Army, it is not feasible or even possible, as it is in the commercial world, to drive and generate detailed maps of all potential routes a priori for autonomous systems. Future military missions will require not only single agents but teams of autonomous vehicles capable of collaboratively determining passable routes over damaged roads and off-road areas as well as through wooded or dense urban environments where prior training data is not available. The Army RAS challenge then is to operate in these complex unknown environments, with little or no infrastructure, at a very high operational tempo while also overcoming system level complexity. Army RAS will not only need to move from point A to point B, they will need to perform tasks along the way. Most efforts today only deal with mission complexity for single RAS or collaborative and homogeneous RAS over a single or sequential set of missions or tasks. Future tasking of RAS will require a much more dynamic tasking and control approach where N multiple agents are performing M several asynchronous tasks and where individual robots can adapt roles as the mission success demands. Another axis of complexity that is unique to Army RAS will be the need to operate in environments that include peer adversarial manned and unmanned systems. This will require systems to communicate in jammed or denied environments, individually and collectively perceive and predict adversarial maneuvers and behaviors, and adapt their behaviors and collectively act in tactical ways to exploit things like windows of opportunity to identify and maintain corridors in contested space or swarm and create overmatch situations.

The goal of the Distributed and Collaborative Intelligent Systems and Technology (DCIST) Collaborative Research Alliance (CRA) is to collaborate with the academic community to study the underlying science questions and develop innovative methods and concepts that will address these capability gaps and enable behaviors that are not brittle and preprogrammed but rather adaptive, resilient, and learnt.

In the process, the CRA will address the assumptions, constraints, and axes of complexity to understand and inform future operational concepts and how diverse, multi-agent RAS can collectively sense, infer, reason, plan, and execute in the face of a peer adversary.

2. DCIST Research Plan

The US Army Combat Capabilities Development Command Army Research Laboratory has established an enterprise approach to intelligent systems that couples multi-disciplinary internal research, analysis, and operations with extramural research and collaborative ventures. CRAs are one of the principal contract vehicles that DEVCOM Army Research Laboratory uses in this enterprise approach to focus on the rapid transition of innovative science and technology for Army Modernization. CRAs are Cooperative Agreements awarded to a Consortia of industry and academia. Collaboration is a key element of the CRA model and together the Consortia and the Government work together through an Alliance where each member brings with it a distinctly different approach to research. This approach enables the Alliance to bring together world class research talent and focus it on Army-specific technology objectives for application to Army priorities.

For the DCIST CRA, the Consortium Lead Research Organization (LRO) is the University of Pennsylvania. The University of Pennsylvania is supported by Georgia Tech and the Massachusetts Institute of Technology as Thrust Area Lead Organizations. Representatives from both DEVCOM ARL and the Consortia Lead Research and Thrust Area Lead Organizations make up the Technical Management Group (TMG) of the CRA. The TMG works to continuously steer the vision and quality of research within the program as well as develop the CRA's Biennial Program Plan (BPP). The BPP covers a two-year timeframe and provides a detailed plan of research activities and objectives for the CRA. The BPP also includes the efforts of the Consortia technical subawardee organizations. This report specifically covers activities under the DCIST CRA BPP for Fiscal Year (FY) 2019–2020, which included subawardees from the Massachusetts Institute of Technology (MIT), Georgia Institute of Technology, University of Southern California, the University of California, Berkeley, the University of California, San Diego, the University of Cambridge, and the New York University.

To address the Army challenges and underlying science questions, the DCIST CRA is organized into three research Thrust Areas:

- **T1: Distributed Intelligence:** Perception-action-communication decision and control loops for distributed and collaborative intelligence, inference, and decision-making for future manned-unmanned teams in dynamic and complex environments.
- **T2: Heterogeneous Group Control:** New methods and architectures and realize speed of battle operation and control of large-scale heterogeneous group behaviors and interactions between humans and autonomous agents.

- **T3: Adaptive and Resilient Behaviors:** Robustness and adaptation of heterogeneous teams to realize resilience to failures, loss, and compromised systems and communications in face of dynamic situational awareness, changing environmental conditions, and adversarial behaviors and capabilities.

To drive the research to relevant Army outcomes and to also foster communication across individual principal investigators and tasks, the CRA is further organized along fourteen capabilities that have been identified as enabling capabilities for future RAS in MDO. These are as follows:

RAS capabilities being developed under the T1: Distributed Intelligence:

- Multi-Agent Situational Awareness
- Collaborative Learning and Intelligence
- Adaptation and Learning in Wireless Autonomous Systems
- Hierarchical Abstractions for Planning
- Joint Resource Allocation in Perception Action Communication Loops

RAS capabilities being developed under the T2: Heterogeneous Group Control:

- Hierarchical & Distributed Control for Adversarial Operations;
- Scalable Task Assignment for Heterogeneous Multi Unit Teams;
- Tactical Engagement of Heterogeneous Teams in Complex Environments; and
- Human Interaction with Large Heterogeneous Teams

RAS capabilities being developed under the T3: Adaptive and Resilient Behaviors:

- Adaptive and Resilient Behaviors
- Resilient Situational Awareness
- Wireless Communication Networks for Distributed Collaborative Systems
- Exploiting Heterogeneity for Resilience
- Multi Agent Behaviors in the Presence of Adversaries
- Rapid Adaptation to Large Disturbances

In addition to the Thrust Area research, the program is also doing research in Cross-Disciplinary Experiments (CDEs) that serve to explore and discover interdependencies across the three research Thrust Areas, build coordinated and cumulative experimentation to demonstrate the Art-of-the-Possible, and answer more operational focus research questions to inform RAS concepts and capabilities for future MDO and drive academic-government collaboration and technology transition. Two CDE capabilities that are reported out here in 1) Heterogeneous Multi-agent Situational Awareness and 2) Dynamic Teaming Operations in Contested Environments.

The research within the DCIST CRA is being transitioned into DEVCOM ARL's AI for Maneuver and Mobility (AIMM) and Emerging Overmatch Technologies (EOT) Essential Research Programs (ERPs). It also has links to the Versatile Tactical Power and Propulsion (VICTOR) and Human-Autonomy Teaming (HAT) ERPs. Through these programs, the methods being developed are informing future concepts involving Next Generation Combat Vehicles and Future Vertical Lift modernization priorities and the Army's Autonomy Priority Research Area.

The following sections provide a brief introduction into each Thrust Area, provide a short description of the relevancy and innovative research being explored to address each capability, highlight some of the key technical accomplishments that have happened during the two-year BPP reporting period, and summarize the potential impact of the findings to date.

3. Methods and Findings: Progress to Date and Future Plans

This section describes the three major research thrust areas and the CDEs that focus on developing capabilities to address scenarios described earlier. Each thrust summarizes the technical objective, list of capabilities planned, major accomplishments to date, and the future plans.

3.1 Thrust 1: Distributed Intelligence

A large team composed of humans, ground vehicles, and aerial vehicles with various prior information and sensing, memory, and computation capabilities should be able to plan and execute complex missions in an uncertain, contested, dynamic environment in a way that is more effective, robust, and resilient than any individual agent. Robust and resilient collective behavior requires the ability to maintain shared situational awareness in dynamic environments with different spatio-temporal scales, achieve joint tasks whose specifications involve unknown parameters such as identities of people or locations of objects, and do so while operating with degraded and disrupted network communications.

The conventional approach to providing such a shared representation across multiple agents is to design the system with a central processing unit (a base station or “the cloud”) that maintains a consistent model of the world at a fixed level of representation (e.g., fixed spatial resolution) that is updated in real time as information is received from other agents. While this approach may successfully achieve robust performance in small-to-medium teams, it has many drawbacks: the central unit creates a single point of failure; the system cannot respond swiftly to new data; an all-to-one communication topology cannot guarantee satisfactory levels of quality of service to all agents in a large team; and, if agents experience loss in communication to the central unit, they are forced to retain all gathered information until they reconnect. Systems that deviate from all-to-one communication topologies are typically based on hand-coded abstractions of sensors and actuators and consensus algorithms that attempt to achieve the same representation across all agents. These hand-coded abstractions are brittle in the face of dynamic environments, especially in adversarial settings, and frequently do not scale to large numbers of agents.

Recent advances in sensing, control, and autonomy allow a single agent to build an environmental model from sensor data and generate plans through that environment, and may even be able to design point solutions for centralized multi-agent systems. However, basic research is needed to develop the foundational principles for designing a truly distributed system that can accomplish the following:

- Build and distribute hierarchical and composable representations of heterogeneous information;
- Perform general data fusion, inference, and planning with these distributed representations;
- Incorporate learning and adaptation to enable these representations to flexibly handle system heterogeneity and dynamics; and
- Move beyond assumptions of network availability and permissiveness and perform sensing, acting, and learning that is cognizant of constraints in communication resources.

These challenges are addressed by driving and coordinating Thrust 1 technical research toward five enabling capabilities:

- Multi-Agent Situational Awareness,
- Collaborative Learning and Intelligence,
- Adaptation and Learning in Wireless Autonomous Systems,

- Hierarchical Abstractions for Planning, and
- Joint Resource Allocation in Perception-Action-Communication Loops.

3.1.1 Capabilities Description and Major Accomplishments

3.1.1.1 Multi-Agent Situational Awareness

Advanced situational awareness in complex, unstructured, and dynamical environments require autonomous agents in multi-agent teams to be able to build and maintain multifaceted models of the environment, including geometric abstractions (useful for safe navigation), semantic labeling (useful for identification and characterization of mission-relevant entities), and temporal dynamics (useful for op-tempo situational awareness). It is of great interest to develop foundational tools and theory supporting efficient spatiotemporal fusing of distributed metric-semantic information across a heterogeneous team, adapting to the computation, information, and communication resources of individual agents to achieve global situational awareness. Some underlying research topics that need to be addressed include the following:

- Representations and algorithms for joint and hierarchical modeling of geometry, semantics, and physics in real time.
- Object shape, appearance, and dynamics models as well as hierarchical and composable abstractions.
- Learning and inference algorithms that utilize data available from prior experience as well as online real-time observations.
- Algorithms that quantify uncertainty and provide probabilistic guarantees, which are critical for decentralizing the representations and enabling adaptive behaviors.

This capability is developing new representations and algorithms that are able to perform efficient probabilistic inference over disparate forms of information such as spatial, semantic, temporal, and topological. The capability will also develop resource-aware techniques for deciding asynchronously how and when to communicate data and models to teammates and to fuse and update the environment representations. Our research agenda is investigating online object-level mapping; distributed Bayesian inference; construction of dynamic scene graphs; association, alignment, and fusion of incomplete and heterogeneous representations; distributed pose graph optimization; and uncertainty quantification and propagation in dense metric-semantic environment abstractions.

Research under this capability has developed several significant underlying software approaches for enabling multi-agent representations of geometric and semantic information for enhanced localization, mapping, planning, and situational awareness. One of them, Kimera-Multi (Chang 2020), is the first system for distributed and dense metric-semantic simultaneous localization and mapping (SLAM) that can use a team of robots to gain situational awareness over a large environment, under realistic constraints on communication bandwidth, local sensing, and computation at each robot. The goal is to estimate a metric-semantic 3-D model of the environment that describes the geometry of the scene the robots operate in (e.g., presence and shape of obstacles), as well as its semantics, where the robots are tasked with annotating the scene with human-understandable labels in a given dictionary (e.g., “building,” “road”).

Kimera-Multi brings together key distributed localization and mapping capabilities as each robot runs single-robot Kimera (including Kimera-VIO and Kimera-Semantics) to estimate the local trajectory and a mesh representation of the environment. Robots then communicate to perform distributed loop closure detection and outlier rejection. Globally consistent estimation is achieved using the novel distributed pose graph optimization (PGO) algorithm developed in this task (Tian 2020), which has been shown to outperform state-of-the-art distributed PGO methods. Kimera-Multi has been evaluated in two large-scale simulation scenarios using the DCIST unity simulator and three ground robots (see Fig. 1). Results show that Kimera-Multi is an efficient, accurate, and robust solution for distributed metric-semantic SLAM and 1) is able to build accurate 3-D metric-semantic meshes, 2) is robust to incorrect loop closures while requiring less computation than state-of-the-art distributed SLAM backends, and 3) is efficient, both in terms of computation at each robot as well as communication bandwidth.



Fig. 1 Dense metric-semantic 3-D mesh model of an urban scene generated by Kimera-Multi with three robots

A second accomplishment was the development of the Consistent Lifting, Embedding, and Alignment Rectification (CLEAR) framework (Fathian 2019) for fusing observations across agents to obtain a common representation. The conventional approach to building such shared representations is to fuse observations incrementally between agent pairs, which can yield inconsistent representations with low accuracy. The multi-way framework of CLEAR improves alignment and fusion accuracy by associating observations jointly across all agents, and the research addressed technological gaps in previous methods, such as high computational complexity and inconsistencies. CLEAR was benchmarked on both synthetic and real-world data sets; showcased in the DCIST photo-realistic simulator to fuse object-level abstractions of the environment across multiple robots; and evaluated experimentally in a collaborative mapping mission to construct a fused occupancy map of a forest canopy (see Figs. 2 and 3). These evaluations confirmed that CLEAR has 1) several orders of magnitude performance improvement over state-of-the-art techniques in both accuracy and runtime, 2) superior scalability to large-size problems, and 3) low computational complexity that is essential for real-time missions. Continuing research is focused on 1) extending the framework to fuse heterogeneous representations obtained from sensors with different modalities, 2) integrating geometric outlier rejection techniques to extend framework’s resiliency to large noise and outliers, and 3) developing methods of incorporating out-of-order, asynchronous data in the fusion process to account for delay and packet loss in challenging communication regimes.



Fig. 2 An aerial vehicle in a collaborative forest mapping mission




Fig. 3 Fused 3-D occupancy grid of environment and vehicle trajectories at the end of the mission. Vehicle paths are shown in blue and yellow and each coordinate frame represents the origin of a submap.

More details for the research on Kimera-Multi and CLEAR can be found in references (Fathian 2019; Chang 2020; Rosinol 2020a; Tian 2020a, 2020b). In addition to Kimera-Multi (Chang 2020) and CLEAR (Fathian 2019), this capability has advanced algorithms for single-platform object-based SLAM (OrcVIO) (Shan 2020), robust data association (CLIPPER) (Lusk 2020), and Distributed PGO (Tian 2020), additional research developments under this capability can be found in Atanasov (2018), Feng (2019, 2020), Paritosh (2019, 2020), Rosinol (2019, 2020b, 2020c), Chamon (2020), Duong (2020), Lajoie (2020), Milano (2020), Shan (2020), and Zobeidi (2020). Together, this research has created a framework that will enable a distributed multi-robot spatial perception engine that builds metric, semantic, and/or abstracted models of the environment in real-time and under limited communication.

Impact: The impact of this research is that, during a maneuver operation by multiple autonomous systems, the autonomous systems will use this software to create and align geometric and semantic information to reduce uncertainty and enable better individual and team localization, mapping, and path planning in complex dynamic environments.

3.1.1.2 Collaborative Learning and Intelligence

The key developments are focused around learning behavioral skills suitable for multi-agent coordination, and developing constrained optimization and learning methods that can utilize such skills to solve complex multi-agent coordination problems. There are two key research highlights in this capability that both relate to *new learning systems*. The first contribution consists of a framework to train a *vision-based navigation system* entirely using autonomously collected data. The second contribution consists of a new model for learning to *coordinate multi-agent systems where individual agents have distinct objectives*. These two contributions are detailed next.

We developed a framework for vision-based robotic navigation that can readily generalize to diverse real-world environments, operate in GPS-denied environments, and enable simple and accessible goal specification by humans or upstream planning algorithms through desired locations, images of goal landmarks, and latent (learned) goal representations. The core concept behind our framework is to *train a vision-based navigation system entirely using autonomously collected data*: all data collected by one or more mobile robots is collected into a single large data set, and then used to learn models of reachability, traversability, collision avoidance, and other navigational affordances. In a series of papers, we present the capabilities of this framework. The BADGR (Kahn 2021a) system uses unlabeled navigational data to learn how to reach user-specified destinations while satisfying desired constraints, such as staying on paved roads, avoiding paved roads, or avoiding collisions. This system can navigate through urban environments, as well as off-road environments including tall grass, where the learned model can determine what types of vegetation are traversable and what types are impassable obstacles. The follow-up LaND (Kahn 2021b) system incorporates online user feedback (in the form of safety aborts) to further learn how to navigate urban environments while following user-desired conventions, such as staying on sidewalks. This system was able to navigate several kilometers of Berkeley sidewalks autonomously, using only onboard camera observations (without LIDAR, GPS, or other sensors). The ViNG (Shah 2021a) system (as illustrated in Fig. 4) further extends this capability to reach distant goals specified by goal images: a human user can take a photograph of a desired destination, and the robot will navigate to this destination by using the learned model and a topological graph representation of landmarks. The RECON (Shah 2021b) system further extends this capability to search for visually indicated targets in new environments, performing a fringe-based search using the same learned model. All of these systems share a common backbone and can all be trained on all data collected by the mobile robot platforms, without any human labeling (with the exception of LaND, which also uses human-provided labels to learn the semantics of navigating sidewalks in cities).




Fig. 4 Learning robotic navigation: ViNG builds and plans over a learned topological graph consisting of previously seen egocentric images, and uses a learned controller to execute the path to a visually indicated goal. Unlike prior work, our method uses purely offline experience and does not require a simulator or online data collection. Note that the graph constructed by our algorithm is not geometric and nodes are not associated with coordinates in the world, but only with image observations—the top-down satellite image is provided only for visualization and is not available to our method.

When considering multi-agent systems with multiple, distinct objectives, the standard design choice is to model these as separate learning systems. Such a design choice, however, precludes the existence of a single, differentiable communication channel, and consequently prohibits the learning of inter-agent communication strategies. We addressed this gap by presenting a learning model that accommodates individual non-shared rewards and a differentiable communication channel that is common among all agents. We developed a new model for learning to communicate to coordinate multi-agent systems in the presence of agents with potentially conflicting objectives (Blumenkamp 2020). This model consists of three key components, 1) a monolithic, decentralizable neural architecture that accommodates multiple distinct reward functions and a common differentiable communication channel, 2) a reinforcement learning algorithm that elicits the emergence of strategic communications, and 3), a post-hoc interpretability technique that enables the visualization of communicated messages.

The experimental evaluation is based on a multi-agent system with a mix of cooperative and self-interested agents and demonstrates the effectiveness of the learning scheme in multi-agent coverage and path planning problems (Fig. 5). Results show that it is possible to learn highly effective communication strategies

capable of manipulating other agents to behave in such a way that it benefits the self-interested agents. Overall, we demonstrate that adversarial communication emerges when local rewards are drawn from a finite pool, or when resources are in contention. We also show that self-interested agents that communicate manipulatively, however, need not be adversarial by design; they are simply programmed to disregard other agents' rewards.




Fig. 5 Cooperative multi-agent tasks: Overview of grid-world environments used in our experiments. Cooperative and self-interested agents are visualized as blue and red squares, respectively. Black cells correspond to obstacles. In the coverage environments, different colors indicate the coverage achieved by individual agents. In the path planning environment, labeled goal locations are indicated by circles.

Impact: The end result are methods that learn in a way that generalizes to multiple tasks and objectives and can learn skills and capabilities in the context and constraints of a team to solve heterogeneous multi-agent coordination problems. The impact of this research is that large heterogeneous robot teams can be trained using small data sets and an optimized centralized solution. This learned heuristic is then generalized so that in practice, such as during a coverage and tracking mission, the team can adapt to previously unseen coverage scenarios, handle communication link losses, address partially known environments, and scale to larger numbers of agents and larger maps than possible in the training phase.

3.1.1.3 Adaptation and Learning in Wireless Autonomous Systems

In addition to sensing and navigating, autonomous agents must work together to solve inference, learning, and control problems. This is the computing that turns the sensor measurements into actionable decisions. Solving these problems optimally requires agents to communicate with one another, as they see only local, overlapping views of the field of interest. In realistic scenarios, the agents are immersed in different environments at different times, are subject to different dynamics and constraints, and their ability to communicate with one another changes rapidly. In addition, decisions must be made in real-time and with limited computational requirements, and the agents deployed might have different compute capabilities. The goal of this capability is to build new mathematical and

algorithmic frameworks for distributed multi-agent learning problems that respects these constraints and to develop the following:

- An understanding for the role that communication plays in our ability to learn models, and the effect of communication limits on decentralized learning, and
- Machine learning frameworks that can be implemented in a distributed manner while preserving the privacy and security of participating agents.

The approach is focused on distributed reinforcement learning for navigation and control as well as federated learning, but many of the general techniques developed might be applied to other inference problems.

Our focus in this subtask over the last year has been on formalizing the multi-task reinforcement learning (MTRL) problem, and understanding when and how it can be solved. As we describe below, we have investigated the optimality conditions for the MTRL problem in the general setting where the tasks may have different state spaces and transition probabilities, and have studied the convergence rates both of a decentralized version of policy gradient and decentralized Q-learning for solving multi-task multi-agent reinforcement learning (RL). The analysis of the algorithms is complemented with experiments on real world MTRL problems.

It has been shown that in single-task RL, there always exists a deterministic optimal policy (under mild assumptions). We have found a relatively simple example demonstrating that the same is not true for the MTRL problem: a deterministic optimal policy may not exist. As many iterative algorithms for RL (e.g., Q-learning) implicitly rely on an optimal deterministic policy, this immediately demonstrates that the multi-agent and single-agent cases are fundamentally different. Another difference is that the objective function for a single-task RL problem is known to obey the so-called “gradient domination” condition, which essentially implies that every stationary point of the objective function is globally optimal (Fig. 6, left). A consequence of this property is that it allows gradient descent algorithms to find the globally optimal solution. We have shown, however, that in general MTRL problems do not enjoy this property, and hence there may exist suboptimal stationary points (Fig. 6, right).




Fig. 6 Optimization landscape of single-task/multi-task RL. Unlike single-task learning problems, multi-task RL problems can have isolated local minima, meaning that there are no a priori guarantees that descent algorithms will find globally optimal solutions.

We have also studied the convergence rate of a decentralized multi-task policy gradient algorithm. We show that in the general setting, the algorithm converges to a stationary point with rate $O(\frac{1}{K})$ (the norm of the gradient is inversely proportional to the number of iterations) and under a further assumption that guarantees the states common between two tasks to be visited equally often, the algorithm finds a globally optimal policy with rate $O(\frac{1}{\sqrt{K}})$. Empirically we test the algorithm on a 4-task simulated drone navigation problem. With the performance measured by the mean safe flight (MSF), the multi-task algorithm learns a common policy effective for all tasks without compromising its performance in any task; this is shown in Fig. 7.




Fig. 7 MSF for (simulated) aerial drones trained to navigate in different indoor environments. In the table on the right SA- i is a drone trained only in environment i ; the policy learned performs well in environment i but poorly in the other three environments. When a single policy is learned for all four environments using distributed optimization (row DCPG), the results are superior across all cases.

Finally, we have studied a decentralized stochastic approximation framework and shown its application to decentralized multi-task Q learning. Under mild regularity assumption, we show that the algorithm converges to a ball around the optimal solution (i.e., the root of a strongly monotone operator) linearly under a constant

step size and converges exactly to the optimal solution with rate $O(\frac{1}{K})$ under a properly selected diminishing step size.

We have continued our study of distributed optimization problems whose structure is described by a graph. The mathematical goal is to solve an optimization problem of the form

$$\sum_i^T f_i(x_{[i]}), \underline{x}_T = (x_0^T, x_1^T, \dots, x_m^T), x_i \in R^{n_i}$$

where each $x_{[i]}$ is a partial selection from the entire subset of variables. We can associate each function f_i with a node on a graph, and draw an edge between nodes if two functions share at least one variable. The key questions we are interested in include: Is there a "fast-updating" scheme with guarantees to update the solution? When adding a node, do we really have to update the entire solution, or can we perform a "limited memory" local update?

Problems of this sort arise in many interesting problems in signal processing, machine learning, control, and statistical inference. For example, in localization and tracking the $x_{[i]}$ are locations/orientations, and f_i capture dynamics and measurements. In multi-task learning, $x_{[i]}$ are decision variables for different tasks, and the f_i are loss functions and some local regularizers.

We began our study with chain graphs (i.e., path); they have a simpler topology and correspond to key problems as tracking and time-varying optimization. Our main theoretical result stems from the observation that under certain mild conditions on the coupling of the loss functions, the update to the solution as we add a new function at the end of the chain decays exponentially as the update travels backward in the chain. In turn, this tells us that the estimate $x_{(T)}$ eventually converges to limit point x_t^* as $T \rightarrow \infty$, and that this convergence is exponential such that $\|x_t^* - x_{t|T}\|_2 \leq \text{Const.} a^{T-t}$ for some positive constant $a < 1$. Capitalizing on the exponential convergence result, we also derived an online-like Newton algorithm that solves such optimization programs while using only finite memory with provable accuracy guarantees.

In addition, we have considered how these results can be extended into more general graphs. As a result, we were able to derive a graph reduction transformation that reduces any graph into a chain graph. The transformation is equivalent in the sense that any theoretical results that can be obtained from the reduced graph also apply to the unreduced graph. This idea is illustrated in Fig. 8. As a preliminary result, we get that, again, local modification to the graph leads to updates that exponentially decay as they travel inside the graphs. We are now working on finalizing these results and an efficient algorithm for computing the updates.




Fig. 8 (a) $-G(V,E)$ and its evolution as new nodes and edges are added—the goal is to update the solution to the associated optimization program efficiently. (b) We reduce the general graph (left) to the chain graph (right); starting from the new node f_8 , expanding one edge outward each time, we collect all unassigned neighbors into a single node.

Impact: The key outcome of this research is a mathematical and algorithmic framework for distributed problems that respects communication constraints. Autonomous agents are expected to work together to solve inference, learning, and control problems. Solving these problems optimally requires agents to communicate with one another where that ability can change rapidly and decisions must be made in real-time with limited or varied computational capabilities. Methods developed here will enable software developers to assess potential performance gains, understand trade-offs, and provide performance guarantees for new classes of learning algorithms (e.g., agent teams that can assess realistic bounds on their joint localization and mapping accuracy as communications degrade).

3.1.1.4 Hierarchical Abstractions for Planning

Future multi-agent autonomous operations will require instantaneous real-time decisions in a rapidly changing environment. The computational complexity of planning in a highly dynamic multi-agent environment grows exponentially with the time and length scales and the number of cooperative and non-cooperative agents in the environment. This specific capability addresses the computational complexity of the planning process using hierarchical planning with abstractions. Some underlying research topics that need to be addressed include the following:

- Design planning models that are hierarchical and composable and that represent the prior and acquired knowledge. Hierarchical representations imply a notion of abstraction, the ability to group similar items together.
- Identify which parts of the world model to incorporate into the planning model, which parts of the world model can be grouped together, and which can be abstracted away without loss of required details.

- An additional research issue is how the abstracted model changes with actions, and predicts performance with respect to the objective function.

The innovative approach taken defines novel complex mission languages to specify mission requirements that can incorporate both geometric and semantic environmental uncertainty, as well as scalable, distributed planning methods that can quickly adapt to continuously learn semantic-geometric models of the environment.

Object-level semantic information can provide important contextual cues beyond the range of dense geometric information to inform more intelligent long-horizon decisions. Despite their intuitive usefulness, object-level maps can be difficult to integrate within planners designed for unknown environments. In this work, we combined the computational power of randomized motion planners with higher-level semantic information via a learned sampling distribution, enabling intelligent navigation in structured, unknown environments. We developed a planning approach that optimizes a predictive sampling distribution inferred from dense geometric representations that track unobserved space, explicit object-level contextual cues both within and beyond the range of dense geometry, and information about the goal.

Plans generated by the Learned Sampling Distribution (LSD) and baseline were compared, demonstrating that a probabilistic road map using our LSD sampler was more likely to find feasible plans than using a sampler informed only by geometry. We note that the learned distribution outperformed the baseline most notably at low sample counts, indicating that our method is especially useful for resource-constrained platforms.

For “Information-Theoretic Resource-Aware Perception and Planning,” we have developed a framework for path-planning on abstractions that are not provided to the agent a priori but instead emerge as a function of the available computational resources. We show how a path-planning problem in an environment can be systematically approximated by solving a sequence of easier-to-solve problems on abstractions of the original space. Specifically, we consider complexity reduction in path-planning problems by means of graph abstractions for resource-limited agents by combining aspects from both the planning and bounded-rational decision-making communities. Our contribution is two-fold. Firstly, we employ an information-theoretic approach to generate multi-resolution abstractions that are not provided a priori for the purposes of path-planning and secondly, our framework couples the environment resolution to the resulting path quality. To the best of our knowledge, there are no existing approaches that utilize information-theoretic abstractions for complexity reduction in path-planning that also guarantee

the monotonic improvement of the path-cost as a function of environment resolution. Coupling the path-cost with the environment resolution provides a link between the path quality, the complexity of executing graph-search algorithms and the information-processing capabilities of the agent determined by the information contained in the generated abstractions.

The example in Fig. 9 shows the utility of the approach and corroborates the theoretical findings. Figure 9a shows the original environment and Fig 9b shows the finest resolution (maximum complexity) path. Example abstract paths are shown in Figs. 9c and d. The figure of merit is the reduction in representation complexity, while preserving the ability to accurately estimate the costs of paths. In this example, utilizing a representation with approximately 70% of the nodes results on average in an abstract path for which the cost is within 30% of the cost of the finest resolution path.




Fig. 9 Planning results for various abstraction levels

In a second sub-task on “Semantic Swarms: Distributed Semantic Planning for Multi-Robot Teams in Unknown Semantic Worlds,” we addressed a motion planning problem for a team of mobile sensing robots with known dynamics that operate in environments with metric and semantic uncertainty. Specifically, the uncertain environment is modeled by static landmarks with uncertain classes (semantic uncertainty) located at uncertain positions (metric uncertainty) giving rise to an uncertain semantic map. The semantic map is determined by Gaussian and arbitrary discrete distributions over the positions and labels of the landmarks, respectively. Such maps can be initially user-specified or can be obtained and updated by recently proposed semantic SLAM algorithms. We considered robots equipped with noisy sensors that are tasked with accomplishing collaborative tasks

captured by a global temporal logic formula in the presence of metric and semantic uncertainty. To account for sensing and environmental uncertainty, we extended Linear Temporal Logic (LTL) by including sensor-based predicates. This allowed us to incorporate uncertainty and probabilistic satisfaction requirements directly into the mission specification. First, we formulated the planning problem as an optimal control problem that designs open-loop sensor-based control policies that satisfy the assigned specification. To solve this problem, we developed a new sampling-based approach that explores the robot motion space, the metric uncertainty space, and an automaton space corresponding to the assigned mission. The open loop control policies can be updated online at time instants determined by an automaton to adapt to the semantic map that is continuously learned using existing semantic SLAM algorithms. To ensure that the proposed sampling-based approach can efficiently explore this large joint space, we built upon our previous works to design sampling strategies that bias explorations toward regions that are expected to be informative and contribute to satisfaction of the assigned specification. We have shown that the proposed sampling-based algorithm is probabilistically complete and asymptotically optimal under Gaussian and linearity assumptions in the sensor models, and shown extensions of the proposed algorithm to account for mobile landmarks and nonlinear sensor models.

Figure 10 illustrates an example problem where a team of $N = 20$ robots need to accomplish a sequence of collaborative subtasks in an environment with $M = 15$ targets, 12 of which are dynamic and 3 of which are static. Some subtasks can be performed in parallel, and have complex temporal constraints between them. The figure of merit is the run-time required to find a sequence of control inputs that satisfies the mission as specified by an LTL formula, and our algorithm is able to generate a sequence of control inputs with terminal horizon $H = 421$ in 14.2 min (Liu 2020; Larsson 2021).




Fig. 10 Robot trajectories that satisfy the LTL task. The robots are initially located at the top left corner. The green and red squares represent the initial and final locations of the targets while the black paths correspond to the target trajectories.

Impact: The key outcomes of this capability are planning models that are hierarchical and composable and that represent both prior and acquired knowledge to deal with complexity and allow more efficient planning. For example, as a single agent or a teams of autonomous agents maneuver through partially known environments, obstacles are revealed during the plan execution. The methods developed here will address the resulting need for fast, on-the-fly, replanning at both the local and global scale for teams of multiple agents at operational tempo and DCIST-relevant length-scales.

3.1.1.5 Joint Resource Allocation in Perception-Action-Communication Loops

Large scale collaborative systems have the potential to satisfactorily accomplish tasks such as exploration, search and rescue, and surveillance; tasks that could otherwise be dangerous or expensive. To realize this potential, these systems have to operate effectively in an unsupervised manner. This means relevant controllers need to be decentralized and rely only on communications between the neighboring autonomous agents that compose the system, avoiding the need for centralized computation units or fusion centers that could become a security threat. Some underlying research topics that need to be addressed include the following:

- The use of graph neural networks for learning naturally decentralized controllers for large scale collaborative systems.

- Training methods including imitation learning and unsupervised self-learning, in order to understand the learning efficiency of various alternatives and how these affect the performance of graph neural networks.
- The robustness of graph neural network-based controllers to changes in the communication network, and to external attacks on the system.

The fundamental innovation in this research resides in using graph neural networks to learn controllers for large scale collaborative systems. Graph neural networks are an extension of graph filters, obtained by adding pointwise nonlinearities. The research approaches the problem of studying the performance of graph neural networks on learning decentralized controllers. It looks at the question of robustness. It addresses the application of graph neural networks (GNNs) to tasks that are relevant to DCIST missions, primarily coverage and navigation in dynamic settings.

This research has been very successful, developing the scientific basis for learning in large scale distributed collaborative autonomy. The capability research has led to a body of work that substantiates the use of GNNs as the technology that enables the learning of collaborative policies that can be implemented in a distributed manner in large scale multi-agent autonomous systems. The fundamental conclusions of this work are 1) the implementation of GNNs can be matched to the wireless communication restrictions of a distributed system; 2) GNNs can be transferred from one system to another because of their transference and stability properties; 3) GNNs can scale solutions that are learned for small numbers of agents to implementations that involve large numbers of agents; and 4) GNNs can leverage centralized protocols to learn distributed approximations. Specific accomplishments are detailed next.

The challenge of collaborative robot teams lies in our need to address the team's Perception-Action-Communication (PAC) loop. We have made significant strides toward solving the integrated PAC loop problem by developing the underlying theory to apply GNNs to multi-agent systems and by addressing challenging Army-relevant tasks such as perimeter defense (Shishika 2020a). We consider a team of defenders that move along a perimeter and must block intruders from breaching (Fig. 11). Critically, we are interested in obtaining a distributed policy that obeys the partial information structure of the problem. Our key insight is that agent-intruder observations and agent-agent communication links may be viewed as edge sets of a multigraph, and local aggregations of information on these graphs obey the information structure. In addition, this graph input naturally addresses the dynamically changing number of defenders and intruders present as the game evolves. In practice, we learn to exchange latent feature vectors as fixed width

messages that efficiently encode task-relevant information and exclude redundant information. The framework is capable of taking advantage of what communication opportunities exist: we recover nearly the performance of an omniscient centralized expert when exchanging multi-byte messages, but can still produce meaningful policies when communication is completely cut off (Paulos 2019).




Fig. 11 Perimeter defense: A team of agents defends a perimeter

In a second success story we considered large scale swarms composed of multiple agents that collaborate to accomplish a task. Individual agents must decide on control actions that are conducive to accomplishing a collective task from their local observations and communication with nearby peers. It has long been known that finding optimal controllers in these distributed settings is challenging. This motivates the use of heuristics in general and the use of learned heuristics in particular. It is germane to emphasize that the challenge in decentralized control stems from the local information structure generated by the unavoidable restriction to communicate with nearby agents. Otherwise, selecting optimal control actions by clairvoyant agents that have access to global information is often not difficult. Building on this observation we propose the use of imitation learning to train policies that respect the local information structure of a distributed system while attempting to mimic the global policy of a clairvoyant central agent. When designing multi-agent systems, we must contend with the dimensionality growth of the system as new agents are added. Furthermore, it is unreasonable to assume that the network during training is the same as the network during execution. Both of these problems can be overcome if we use GNNs. We examine flocking tasks to highlight the ability of our approach to handle dynamic communication networks; see Fig. 12. We show that a global controller outperforms such local controllers, but global approaches are not practical for real deployments. The novelty of our approach to flocking is to aggregate from multi-hop neighbors; this ability allows us to approach the performance of global solutions while respecting realistic communication constraints. Under standard configuration parameters collisions and velocity variances of GNN controllers are within $4\times$ and $7\times$ of the respective metrics of existing decentralized controllers (Tolstaya 2020). Prior to this work,

there has been no principled approach for augmenting the communication between neighbors to pass on information aggregated from multi-hop neighbors.




Fig. 12 Flocking tasks to highlight the ability of our approach to handle dynamic communication networks

In a third successful case study, we focused on the problem of coverage, in which a robot team must visit a set of locations in an environment. We encode the task as a graph: the known map locations and team members are graph nodes, and allowed moves are graph edges; see Fig. 13. A moderate-size coverage task with dozens of goals and fewer than 10 agents can be solved with existing approaches when posed as a vehicle routing problem (Tolstaya 2021). We collect a data set of trajectories generated using the centralized expert solution and use behavior cloning to train a GNN controller to imitate the expert solution. This learned heuristic can then generalize to previously unseen coverage scenarios with more agents and larger maps. The trained GNN models effectively generalize to larger robot teams and map sizes. The models were first trained on 4 agents and 228 waypoints on average. Then, the models were tested on a map size of 5659 waypoints with a graph diameter of 205. The team size varied from 10 to 100 agents. For both the coverage and exploration generalization experiments, the map and team sizes made the centralized expert solution intractable.




Fig. 13 Encoding the task as a graph: the known map locations and team members are graph nodes, and allowed moves are graph edges

Finally, we developed a learning method to solve the unlabeled motion problem with motion constraints and space constraints in 2-D space for a large number of robots. To solve the problem of arbitrary dynamics and constraints we propose formulating the problem as a multi-agent problem. In contrast to previous works that propose using learning solutions for unlabeled motion planning with constraints, we are able to demonstrate the scalability of our methods for a large number of robots; see Fig. 14. The curse of dimensionality one encounters when working with a large number of robots is mitigated by employing a GNN to parametrize policies for the robots. The GNN reduces the dimensionality of the problem by learning filters that aggregate information among robots locally, similar to how a convolutional neural network is able to learn local features in an image. The key takeaway from our experiments is that decentralized inference using this methodology performs always within a small margin (approximately 12–15 s) of the optimal solution and this margin remains more or less constant even if the goals are further away and if the number of robots is increased (Khan 2020). Thus, from this we empirically demonstrate that GNNs trade some measure of optimality in exchange for decentralized behavior and this tradeoff remains more or less constant even as the number of robots are increased (Khan 2019).




Fig. 14 Demonstrating the scalability of our methods for a large number of robots

Impact: The key outcome from this capability is the advancement of GNNs as the fundamental enabling technology for obtaining effective controllers for large scale collaborative systems. A large heterogeneous robot team must visit a set of locations in coverage and tracking task. They are trained using small data sets and an optimized centralized solution; this learned heuristic is then generalized so that in practice the team can adapt to previously unseen coverage scenarios, handle communication link losses, address partially known environments, and scale to larger numbers of agents and larger maps than possible in the training phase.

3.2 Thrust 2: Heterogeneous Group Control

Thrust 2 develops the fundamental understanding and formalism necessary to incorporate heterogeneous robots within operational and tactical planning and execution of mixed Soldier/robot teams. It postulates those theories for developing complex/adaptable control structures within an adversarial setting must be developed to achieve effective alignment of multi-agent team action with strategic goals. It also hypothesizes that a greater diversity of agent capabilities and team compositions (i.e., combined arms formations with associated doctrine) will be required to successfully fight in this new reality, and that the primary challenge for robot teams is to function effectively as part of closely coordinated teams despite the adversary and the complexity of the battlefield.

In order to unlock human–robot collaboration and deliver on the promise of large-scale heterogeneous systems, we must overcome technical challenges in three domains: models, architectures, and algorithms. There is, indeed, as of yet no unifying, formal mathematical modeling framework that fully incorporates sensing, actuation, computation, and communication diversity. As such, existing approaches cannot answer questions about human-robot teaming (HRT) arrangements in general as they rely on ad-hoc, brute force techniques such as explicit enumeration of the differing capabilities of each agent. Among other drawbacks, such a brute-force approach precludes proper analysis of the interplay between task requirements and agent capabilities as well as an understanding of the interplay between the heterogeneous capabilities of complex human and nonhuman agents and groups.

To address these challenges, this thrust is organized around two main pillars, namely *Operational Control* and *Tactical Maneuvers*, driving and coordinating T2 technical research toward four enabling capabilities in Hierarchical & Distributed Control for Adversarial Operations; Scalable Task Assignment for Heterogeneous Multi-Unit Teams; Tactical Engagement of Heterogeneous Teams in Complex Environments; and Human Interaction with Large Heterogeneous Teams.

Operational Control is focused on developing control structures/frameworks that promote robust and adaptive large-scale heterogeneous group control. Specifically pursuing advancements enabling two capabilities:

- Hierarchical & Distributed Control for Adversarial Operations
- Scalable Task Assignment for Heterogeneous Multi-Unit Teams

Tactical Maneuvers is focused on effective tactical (combined arms) maneuvers result in achieving positions of advantage over an adversary and necessitating the effective fusion of movement with an understanding of the battlefield and the

adversary. Effective combined arms maneuvers require tight coordination between heterogeneous agents performing a diversity of roles. Specifically pursuing advancements enabling two additional capabilities:

- Tactical Engagement of Heterogeneous Teams in Complex Environments
- Human Interaction with Large Heterogeneous Teams

The culminating outcome of T2 will be the development of the technological understanding necessary to realize truly collaborative performance in the face of an adversary, while in complex environments, through the optimal exploitation of available agents.

3.2.1 Capabilities Description and Major Accomplishments

3.2.1.1 Hierarchical & Distributed Control for Adversarial Operations

Dynamic engagements against swarming adversaries in realistic battlefield scenarios present a multitude of challenges for the current state of the art technologies and algorithms. For instance, deploying resources (forces, robots, sensors, or supplies) to appropriate locations at the appropriate time is a fundamental problem related to resilient situational awareness, perimeter defense, and many other DCIST-relevant scenarios. This will involve a degree of motion planning as well as allocation of heterogeneous resources across the environment requiring a coherent use of appropriate abstractions, discrete optimization algorithms, game theoretic techniques, and geometric considerations.

This capability will develop theory aimed at strategic deployments of resources over large environments in dynamically changing scenarios involving models of adversarial agents and imperfect/delayed communication. As such there are three innovative contributions to this capability:

- 1) Models, abstractions and algorithms for adversarial teams
- 2) Non-equilibrium strategies for decision-making
- 3) Incorporating heterogeneity and diversity into mission and task planning

Adversarial Teams

The DCIST consortium formulated models for blue teams defending high value targets from attacking red teams in a dynamic setting allowing applications of game theory to solve for optimal defense strategies (Shishika 2020a-b). Specifically, we showed how we might be able to design teams of agents given models of the adversary, the environments, and the resources available to different robots and humans. How a team of defenders can optimally defend against intruders

approaching a convex perimeter by intercepting them, first using a one-on-one defense, and then using a two-on-one defensive strategy. Here the optimal strategy is a Nash Equilibrium where neither the blue team nor the red team has an incentive to deviate from the optimal strategy. The general multi-player game was shown to be NP-hard, but solved using a sub-optimal matching strategy. Finally, an imitation learning algorithm was used to develop a team strategy for large groups in which a model-based algorithm for small teams was used to train perception-action-communication loops for large teams engaged in perimeter defense (Shishika 2020a). This body of work is the first systematic approach to design, analysis, and realization of blue team strategies and implementations on both simulation and virtual platforms used to validate the approaches.

Non-Equilibrium Decision-Making Strategies

The DCIST consortium developed (Tsiotras 2021) hierarchical decompositions for multi-player stochastic games as a means to a) remedy the computational complexity of a general stochastic game; and b) capture heterogeneity in a team that includes agents of different rationality levels. In particular, we formalized and designed “non-equilibrium” strategies, which depart from traditional Nash-equilibrium-based approaches using a Markov Decision Process model as the framework to analyze sequential decision-making under uncertainty.

Heterogeneity for Mission and Task Planning

The DCIST consortium developed a methodology for specifying high-level mission requirements and designed abstractions (traits) of individual agents that were used to synthesize first specifications for tasks and then plans for individual agents. We were able to map trait combinations to measures of efficacy and identify the trade-space of cost versus performance. We were able to perform integrated, multi-level optimization for joint coalition/task assignment, local scheduling, and controller synthesis. At a high level, this can be abstracted as the control of the distribution of traits on a graph, while at a lower level it reduces to task scheduling using MILP/MICP techniques to encode convex team constraints while solving problems efficiently.

Impact: The key output of this capability is the development of an initial theory and the algorithms aimed at strategic deployments of agents in large environments in dynamically changing scenarios involving models of adversarial agents and imperfect/delayed communication. For example, as teams of autonomous systems deploy on an intelligence, surveillance, and reconnaissance (ISR) or perimeter defense task, they would use the methods developed under this project to arrive at planned locations and at planned times when confronted with the realities of fast-

moving swarm-versus-swarm engagements and the presence of imperfect information exchange and non-rational players.

3.2.1.2 Scalable Task Assignment for Heterogeneous Multi-Unit Teams

Given a high-level mission specification, such as patrolling and securing an area or performing a search-and-rescue operation, the capabilities needed to carry out that mission must be assembled by, and distributed across, the participating team of heterogeneous agents. For example, certain sensing modalities might be needed to be able to detect targets, different types of mobility, such as ground and air vehicles, might be required to effectively sweep the area, while different computational and communications capabilities and knowledge bases might be necessary to process and relay information back to a base station. Given a set of such requirements, *this Capability focuses on how to assemble and deploy a dynamic, outclassing, heterogeneous team as a function of the mission specifications.*

One particular manifestation of the heterogeneity concept is mixed human-autonomous agent teams. To be able to reason about such teaming arrangements, one needs computational models that accurately capture human characteristics relevant to the present and future concepts of operation, and must incorporate such models in computation team-level coordination schemes. These models must additionally be able to support team-level assignment tasks. Relevant characteristics will include both physical traits (e.g., fatigue, available weapons), as well as cognitive traits (e.g., trained skills, cognitive load).

Due to the computational complexity associated with optimal assignment strategies in dynamic environments, an enabling, key research issue is how to produce *satisficing* (i.e., good enough) solutions, in a computationally feasible, distributed, and adaptive manner. The search for such a computationally feasible, “satisficing” solution is the primary research issue to be investigated under one task in this capability. In particular, dynamic constraints will be phrased as control barrier functions that ensure that satisfying solutions are always available through the so-called forward invariance property—if the system starts satisfying, it stays satisfying even in the face of dynamic changes.

One prime target application in which the developed framework will be tested is “coverage control” (as interpreted broadly) where the teams are composed of individuals with different sensing, mobility, communications, and computational capabilities. The agents are to spread out in order to cover/explore/protect/attack a given area. As such, coverage control constitutes an enabling capability in a number of assignment problems, including reconnaissance and surveillance.

Research under this capability has shown that by allowing for redundant task assignments, where multiple agents may be assigned to the same task, significant robustness to uncertainty in task performance can be achieved (Malencia 2021). In particular, optimization of *average* costs across all tasks as well as the generalization to the *minimization of worst-case costs* are considered as a way of minimizing the maximum damage done. Key to achieving this latter objective is through the interpretation of a *fair* optimization objective.

Take for example the assignment of robots to deliver medical supplies to wounded Soldiers. Optimizing the average wait time of the Soldiers does not prevent an individual Soldier from waiting arbitrarily long for their supplies, as illustrated in Fig. 15. This performance may cause the system to be perceived as ineffective, unfair, and/or untrustworthy.




Fig. 15 Redundant assignment answers the question, which task should receive extra resources? Given an existing non-redundant assignment and respective cost distributions (top right), a single redundant robot, Robot 2, can be assigned to either task. Utilitarian approaches assign Robot 2 to Person 0 because of the higher improvement in cost (Graph A), whereas fair approaches assign Robot 2 to Person 1 because of their higher need (Graph B).

Specifically, a fairness criterion for redundant assignment based on the philosophy of John Rawls’ *Theory of Justice* was defined and subsequently formulated as a *min-max problem* that captures this fairness criterion and the assignment problem constraints (Rawls 1971). Unfortunately, solving this problem fully is strongly NP-hard. But, by leveraging the particular structural properties of the mathematical problem formulation a polynomial-time, near-optimal solution is found through the introduction of a second decision variable that lets the problem be decomposed using half interval search combined with greedy assignment.

Following the theme of computationally feasible, heterogeneous assignment algorithms, an additional result under this task, reported in (Emam 2020), involves the notion of good enough assignment, where the assignment specifications are turned from costs (optimality) to constraints (feasibility). Recent results in this general area show how such a framework can be made robust and take environmental disturbances or unknown phenomena into account. Without such robustifying measures, the quality of both the allocation and the execution of tasks by the robots may be negatively affected. More importantly, even small environmental disturbances may result in the deterioration of the estimated specialization of the robots at performing tasks. In other words, the framework cannot distinguish between disturbances that the robots actually can and cannot overcome. This negatively affects the ability to allocate tasks based on specialization in an effective fashion.

Motivated by these limitations, a novel framework using Gaussian processes and robust control barrier functions is developed as part of this task that lets the robots learn and model the disturbances in conjunction with assurances of the task execution under these disturbances. Which further allows the robots to distinguish between disturbances that the robots can and cannot overcome; the former being due to model errors, while the latter is caused by the actual incapability of the robots at performing the task. This construction is illustrated in Fig. 16.




Fig. 16 By learning, in real-time, models of environmental disturbances and individual agents' suitability to particular tasks, an adaptive and robust framework is obtained for task allocation and execution for heterogeneous teams.

Impact: The key output from this capability is exploring how to produce sufficient or “good enough” solutions, in a computationally feasible, distributed, and adaptive manner. For example, as a team of heterogeneous robots executing multiple tasks (such as perimeter defense, reconnaissance, security, resupply, etc.) experience unexpected terrain changes (affecting mobility), weather events (affecting sensing), and adversarial attacks (affecting sensing and communications) degrading task

performance resulting the agents would use the methods developed under this task to rapidly adapt and retask (not waiting for optimal solutions that may not exist given the complexity and context of the task) to ensure mission success. As such, the measures of success of the capability are given in terms of an order of magnitude increase in the size of the heterogeneous assignment problems that can be solved rapidly (from 10s of robots to 100s of robots) as well as quantifiable robustness to uncertainties associated with the assignment parameters.

3.2.1.3 Tactical Engagement of Heterogeneous Teams in Complex Environments

This capability focuses on the development of abstractions, algorithms, and frameworks to address two challenges with a particular emphasis on heterogeneity, and adversarial behavior: autonomous tactical team behavior and tactical team composition. Perimeter defense is embedded and essential in many Army missions. In future perimeter defense operations, it is envisioned those robotic vehicles will help assist monitoring and neutralizing threats. It is also generally assumed that dynamic team formation (i.e., the formation of teams based on the threat to be addressed) should perform better, when compared to the static teams, agents in which the teams are determined a priori before any information about the threats become available. This task addresses the above questions, resulting in a formal representation that enables the system to effectively characterize human skills, and aid high-level coordination algorithms, such as task assignment and team composition.

Designing controllers and compositions for heterogeneous multi-robot coordination in complex and unknown environments requires considerable expertise and a significant amount of manual effort. To circumvent these demands, end-to-end methods for learning multi-agent policies have been proposed. However, these often require vast amounts of data, are not interpretable, and are limited to low-level coordination. We have developed a structured approach shown in Fig. 17.




Fig. 17 Overview of the proposed two-level approach to learning heterogeneous multi-robot coordination

Our investigation of tactical team behaviors in partial information target defense games are focused on a scenario in which an autonomous defender is tasked with intercepting an intruder that tries to reach a target region (Shishika 2021). Unlike the original target guarding problem and its various extensions, we consider the effect of partial information by imposing sensing limitations on the robots (Fig. 18). A major accomplishment is the spatiotemporal decomposition of the game into three phases: deployment, asymmetric information, and engagement phase. Focusing on a particular parameter regime, we propose a defender strategy together with the lower bound on its probability of win. A surprising outcome of this study is the emergence of a “see-wait-strike” strategy, whereas waiting is generally suboptimal for the full information game. The defender strategy in each phase is constructed so that the subsequent phase starts in a desired initial configuration. The proposed strategy is also robust against parametric uncertainty.




Fig. 18 A limited-sensing model for target guarding

We have begun a theoretical investigation of perimeter defense using a 1-D parameterized line segment (Fig. 19). The defenders on the perimeter are constrained to move on the line with predetermined maximum speed. The adversarial agents appear uniformly randomly and descend on the unit interval with a unit speed until they breach the perimeter. An adversarial attacker agent is captured if a defender is present at the point of breach of the attacker. Given a set of attacker positions and their maximum speed, we compute the minimum number of defenders to neutralize the attackers. We solve this problem by constructing a directed acyclic graph whose nodes represent the attackers. We show that the problem can be solved using the well-known maximum bipartite matching algorithm.




Fig. 19 The red triangles represent the attackers, and the blue circle denotes the defender. At least three defenders are required to neutralize attack.

Impact: The key output from this capability are the methods and tools for the analysis and development of cooperative team strategies in the presence of adversarial agents. For example, a team of heterogeneous agents, using learned coordination policies from observations of an expert team, engage an adversarial team and infer the opponents decision-making process, then decompose the engagement into different phases based on information asymmetry to create windows of opportunity and tactical advantage (Pierpaoli 2020; Silva 2019).

3.2.1.4 Human Interaction with Large Heterogeneous Teams

We envision heterogeneous multi-agent teams, consisting of both autonomous and human agents, coordinating seamlessly to accomplish complex task objectives. Toward this goal, much research has focused on how to coordinate *autonomous* agents within such a team, but significantly less work has examined how *human* operators can work effectively as part of a human–robot team. Research in this task focuses on identifying techniques that enable improved human interaction with, and control of, a heterogeneous robot team. Some underlying research topics that are addressed include:

- Design of formal representations that characterize human skills and model how humans perceive and communicate time-evolving information.
- High-level coordination algorithms, such as task assignment and team composition, that leverage the above representations to demonstrate improved human–robot coordination in complex environments. Algorithms that enable multi-agent task models to be actively learned from human operator input.

Our research represents a multifaceted and multidisciplinary approach to studying the role and interaction dynamics of humans in human–robot teams. Our work has contributed new insights into human perception and communication, including how

to optimally compress communication while maintaining accuracy (Lynn and Bassett 2020). Leveraging these insights, we contributed techniques for measuring and modeling human capabilities and traits relevant to multi-agent teaming (Kolb et al. 2021); such models help improve team performance through more effective task assignment (in preparation). To learn models of multi-agent tasks, we contributed novel learning methods that actively maintain value alignment and model the situational awareness of both robotic and human agents (Shannon et al. 2017). Finally, we seek to not only build a priori models of humans, but also track and respond to dynamic changes in human operator abilities over time. Toward this goal we are working on active techniques for tracking user performance over time (in process).

Human perception and communication: Complex environments are characterized by many events among which the pattern of possible event-to-event transitions has non-trivial topological structure. Each event leads to a specific set of other events, with some probability. The event-to-event path leading from an adversarial interaction (a potentially negative event) to a desired outcome (a positive event) can be long, circuitous, and not always apparent to any given Soldier or multi-agent team on the field. Similarly, the event-to-event path leading from the movement of one robot in the swarm to a conformational change in the swarm signaling a message of import to the human is complex, and must be accurately inferred by the human in order for them to anticipate outcomes and respond swiftly. Human, robot, and human-robot communication regarding event patterns similarly displays a non-trivial topology, as do the set of actions and plans that must be used to effectively engage in the complex environment. Thus, humans and robots are faced with the challenges of accurately perceiving and communicating the topological structure of event-to-event transitions in complex environments, and choosing the appropriate action-to-action transitions in response to that environment. Moreover, they must often solve these challenges under massive constraints of time and limited sensing, leading ultimately to sparse data (Suri et al. 2019).

The study of human perception and communication of event transitions has typically focused on pairs of events: A and B (Fig. 20). This focus has limited the resultant insights to simple environments, elementary agents, and short time horizons. PI Bassett’s lab has pioneered a transformative approach to the study of human perception and communication of event transitions arising in complex environments, performed by sophisticated agents, and evolving over long time horizons. We define graph learning to be the process whereby humans learn and represent the networks of event transitions in the world around them. Using this innovation, we have determined optimal models for the compression of communication while maintaining accuracy (Lynn and Bassett 2020), and have

made significant progress on the control of human perception of the environment (in preparation).




Fig. 20 Principles of design for optimal communication in multi-agent teams under constraints of time and limited sensing

Modeling human traits: HRT enables groups of humans and autonomous robots to communicate, coordinate, and collaborate to perform a joint activity. Human agents are highly heterogeneous with respect to their innate abilities. Indeed, prior work has shown that humans varied up to 87.5% in two traits associated with active robot path planning (Shannon et al. 2017). Such examples of heterogeneity motivate the need for careful allocation of agents to the various tasks that need to be carried out in an HRT setting. Prior work on task allocation in human-robot teams has largely ignored this variation in favor of simpler aggregate models (Ravichandar et al. 2020). In particular, it is often assumed that all human agents within a given category (e.g., Soldier, firefighter, rescuer) have approximately equivalent attributes and can therefore be assigned arbitrarily. Treating all human operators as identical fails to account for individualized differences in capabilities across operators.

In this work, we model natural variations in human operator capabilities, and then study whether these variations translate into differences in individual performance on HRT tasks. Specifically, we introduce two cognitive tests that measure human cognitive capabilities that are pertinent to interacting with and coordinating multiple robots. Our central research question is to determine whether individual performance on certain cognitive tests serves as a predictor for HRT performance. If so, this information can serve to improve heterogeneous task assignment frameworks, such as Stochastic Trait-based Task Assignment (STRATA). Our results (Kolb et al. 2021) demonstrate that our cognitive tests can be used to effectively predict operator performance on certain tasks, but not others (Fig. 21). Specifically, the situational awareness test predicts performance on the remote teleoperation task, and the network test (shown below) predicts performance on the ad-hoc networking task. This finding indicates that targeted cognitive tests can be developed to quickly and effectively probe individual human abilities prior to task assignment.




Fig. 21 The Network Inference pretest (left) is a simple 5-min cognitive test. Human operator performance on this pretest is highly correlated (middle) with operator performance on complex multi-robot teaming tasks, such as teleoperating multiple ground and aerial vehicles while maintaining an ad-hoc communication network (right). The ability to quickly determine which human operators are better suited for robot control can help inform task allocation in complex HRT scenarios.

Impact: The key outcome from this capability is improved modeling techniques for how humans perceive and communicate time-evolving information and frameworks for modeling human capabilities, with the goal of improving human–robot coordination in the context of teams operating in complex environments. As an example, pre-tests are used to predict an individual person’s ability in controlling a robot swarm and adapt communication of time-evolving information with the robot team. Then models and methods that account for individual human capabilities are used in high-level coordination activities such as task assignment and team composition in mixed human–robot teams to improve coordination and performance.

3.3 Thrust 3: Adaptive and Resilient Behaviors

The DCIST vision calls for a team of heterogeneous agents providing increased situational awareness, warfighter standoff, and dilemmas to the adversary. These capabilities, building on the foundations of Distributed Intelligence and Heterogeneous Group Control, must thrive in dynamically evolving and potentially unknown and hazardous environments. Successful systems must encode resilience to disturbances into their design. These disturbances go beyond passive effects due to the environment and extend to include effects driven by an adversary force. Attrition of agents, rapid changes in the environment, and unexpected or deliberate loss of wireless communication all represent the operational conditions under which autonomous technology on the battlefield must operate.

Traditional approaches address resilience primarily with the tool of over-provisioning. Assigning larger numbers of agents to perform a task or additional communication links to a network in order to overcome modeled losses of these resources at runtime. Indeed, there is a rich and growing literature for the theory of resilience even in the challenging space of networked control systems as are found in many modern infrastructures (e.g., the power-distribution grid). However, the Army-specific challenges encountered in MDO may not always allow for gross over-provisioning and overmatch will likely come by operating closer to the edge of the envelope in terms of how assets are allocated. Likewise, changes in the environment or interaction with other entities are often addressed through replanning, adaptive control, or online learning. Unfortunately, many of these approaches can be exploited by an adversary, putting autonomous systems at risk of always operating reactively rather than proactively, which leads to inadequate performance.

Our approaches will strive to maximize robustness at the micro level of individual perception, inference, or learning modules subject to the available resources, and seek resilience at the macro level of the team. Robust perception and learning strategies must be developed to recognize and reject unexpected or adversarial inputs at the micro level, while resilient collaborative control and decision-making strategies must be developed to cope with compromised and failed agents and communication loss at the macro level. The challenge is to design strategies that cope with all of these issues in a general and foundational manner and, yet are practically realizable across various platforms, computing substrates, and communication networks. Our goal in this research thrust is to develop a science of resilience for teams that explores performance/resilience tradeoffs subject to mission context.

Furthermore, in this research thrust, we will utilize a powerful unified representation of uncertainty across many DCIST characteristics (team composition, heterogeneity, environment, context, and communications) and update it rapidly in time, despite its large-scale and distributed nature. In turn, we will develop adaptive group behaviors that are aware of increasing uncertainty in the surroundings as well as models of the adversary and explicitly seek to trade-off uncertainty reduction with task performance optimization.

In order to realize adaptive and resilient multi-agent autonomous systems on the battlefield there are many technical challenges that follow from the research gaps and goals described above. We organize them into two high-level pillars or capabilities: 1) Scalable Multi-Agent Information Acquisition and 2) Adaptive Planning for Adversarial Actions and Large Disturbances. The first is about realizing the vision of improved situational awareness while maintaining increased warfighter standoff and coverage. The second is fundamentally about providing force multiplication, increasing the number of dilemmas presented to an adversary while rapidly adapting to a changing landscape.

Scalable Multi-Agent Information Acquisition is focused on the task of doing resilient multi-agent target detection and tracking. It is spearheaded by efforts to tackle the problems of resilience at scale. As previously described, methods exist to pre-plan for over-provisioning with respect to failure models. Previous DCIST work has spearheaded techniques to tackle these resilience problems dynamically and current efforts seek to increase the scale at which these algorithms can be implemented. Second, situational awareness is constrained by a communication infrastructure that is robust to the failure of individual links or even whole swaths of the network. We aim to address this problem by developing techniques that dynamically allocate wireless network infrastructure in ways that explicitly consider the uncertainty of channels, mobility of the agents, and actions of potential adversaries. Finally, we are studying how to exploit heterogeneity as a means to achieve resilience.

Adaptive Planning for Adversarial Actions and Large Disturbances is focused on working explicitly in the context of the adversary. Rather than blindly reacting to disturbances, this pillar gets at the technical challenges of operating in a fundamentally adversarial environment. Initially this focuses on algorithms for distributed estimation of the adversary state and prediction of intent and future actions. These approaches are challenged by sensing uncertainty and limited knowledge of the environment. Key to this idea is the development of theory and algorithms to explicitly learn how to plan in the context of uncertain knowledge of the adversary. Indeed, online learning in the context of large disturbances is critical

and we aim to push the state of the art in terms of providing guarantees for fast and robust online learning.

These challenges are addressed by driving and coordinating T3 technical research toward five enabling capabilities:

- Resilient Situational Awareness
- Wireless Communication Networks for Distributed Collaborative Systems
- Exploiting Heterogeneity for Resilience
- Multi-Agent Behaviors in the Presence of Adversaries
- Rapid Adaptation to Large Disturbances

3.3.1 Capabilities Description and Major Accomplishments

3.3.1.1 Resilient Situational Awareness

The objective of this capability is to establish the foundations of resilient distributed situational awareness for DCIST. By situational awareness we mean behaviors such as mapping, target identification, localization, and adversarial tracking, in the presence of sensor failures and/or compromised agents. Situational awareness in DCIST scenarios requires the deployment of a mobile team of robots, where each robot needs to be agile; coordinate its motion with its team in a decentralized way; and navigate itself in unknown, complex, and GPS-denied environments, with the objective of gathering the most information about the environment or target of interest. This capability will aim to provide zone reconnaissance for local and global situation awareness using a team of heterogeneous robots in the presence of failures or attacks.

Existing state-of-the-art situational awareness approaches focus mainly on computational efficiency for maximizing information without offering resilience-to-failure mechanisms, especially for distributed or heterogeneous DCIST architectures and behaviors. Furthermore, they are myopic, non-scalable, not-energy efficient, and not resilient to changes or failures to the DCIST robot team. Addressing these challenges requires shifting the perspective of off-line learning from big data and maximizing the performance of every available asset to a parsimonious approach, where only reliable data, sensors, and actions are intelligently selected to obtain sufficient and resilient situational awareness.

The problem of designing the motion of a team of mobile robots to infer the state of an unknown process is known as active information gathering. The first objective of this capability is to generalize active sensing techniques to handle distributed inference and control in multi-robot teams, supporting nonlinear sensing models,

visibility and collision constraints, and adaptation to dynamic changes in the geometry, semantics, or communicability of the surroundings. A second objective is to consider failure-prone and adversarial environments where the robots can get attacked, their communications channels can jam, or their sensors can fail. In such failure-prone or adversarial scenarios, resilient design against worst-case and system-wide failures and attacks becomes important. We developed behaviors for resilient active information gathering that go beyond the traditional objective of uncertainty minimization and guards against worst-case failures or attacks that can cause the withdrawal of robots from the information acquisition task. Resilient active information gathering with mobile robots is a computationally challenging task, since it needs to account for all possible removals of robots from the joint motion design task, which is a problem of combinatorial complexity.

Our research agenda pursued 1) distributed inference and dynamic programming formulations of exploration, active mapping, and target search; 2) distributed submodular optimization; 3) risk-sensitive formulations of estimation, control, and reinforcement learning; 4) value iteration and policy gradient algorithms for active sensing; 5) the development of a software infrastructure of active sensing algorithms for the DCIST consortium. Some highlights include the following:

Energy-aware active information acquisition. A collaboration across DCIST recently considered the problem of planning trajectories for a team of sensor-equipped robots to reduce uncertainty about a dynamical process. Optimizing the trade-off between information gain and energy cost (e.g., control effort, distance traveled) is desirable but leads to a non-monotone objective function in the set of robot trajectories. Therefore, common multi-robot planning algorithms based on techniques such as coordinate descent lose their performance guarantees. Methods based on local search provide performance guarantees for optimizing a non-monotone submodular function, but require access to all robots' trajectories, making it not suitable for distributed execution. We recently proposed a distributed planning approach based on local search that shows how lazy/greedy methods can be adopted to reduce the computation and communication of the approach. We demonstrated the efficacy of the proposed method by coordinating robot teams composed of both ground and aerial vehicles with different sensing/control profiles and evaluated the algorithm's performance in two target tracking scenarios. Compared to the naive distributed execution of local search, our approach saves up to 60% communication and 80–92% computation on average when coordinating up to 10 robots, while outperforming the coordinate descent-based algorithm in achieving a desirable trade-off between sensing and energy cost. This accomplishment, published recently in (Cai et al. 2021), was recently highlighted by MIT News at <https://news.mit.edu/2021/robots-collaborate-search-0513>.

Scalable multi-robot information acquisition. In (Kantaros and Pappas 2021), we recently proposed a novel highly scalable nonmyopic planning algorithm for multi-robot Active Information Acquisition (AIA) tasks. AIA scenarios include target localization and tracking, active SLAM, surveillance, environmental monitoring and others. The objective is to compute control policies for multiple robots that minimize the accumulated uncertainty of a static hidden state over an a priori unknown horizon. The majority of existing AIA approaches are centralized and, therefore, face scaling challenges. To mitigate this issue, as shown in Fig. 22, we proposed an online algorithm that relies on decomposing the AIA task into local tasks via a dynamic space partitioning method. The local subtasks are formulated online and require the robots to switch between exploration and active information gathering roles depending on their functionality in the environment. The switching process is tightly integrated with optimizing information gathering giving rise to a hybrid control approach. We showed that the proposed decomposition-based algorithm is probabilistically complete for homogeneous sensor teams and under linearity and Gaussian assumptions. We provided extensive simulation results showing that the proposed algorithm can address large-scale estimation tasks that are computationally challenging to solve using existing centralized approaches. As demonstrated in (Kantaros and Pappas 2021), this approach allows us to push the scale by an order of magnitude (100+ robots).




Fig. 22 Decomposition of global active information acquisition tasks into local tasks

Reinforcement learning with information theoretic objectives. Information Acquisition for multi-robot systems can be formulated as a multi-agent reinforcement learning problem (MARL). MARL is prone to the straggler effect where some learners are slower than others. Stragglers arise frequently in a distributed learning system, due to the existence of various system disturbances

such as slow-downs or failures of compute nodes and communication bottlenecks. To resolve this issue, in (Wang et al. 2021) we recently proposed a coded distributed learning framework, which speeds up the training of MARL algorithms in the presence of stragglers, while maintaining the same accuracy as the centralized approach. As shown in Fig. 23, a coded distributed version of the multi-agent deep deterministic policy gradient (MADDPG) algorithm is developed and evaluated in (Wang et al. 2021). Different coding schemes, including maximum distance separable (MDS) code, random sparse code, replication-based code, and regular low-density parity check (LDPC) code are also investigated. Simulations in several multi-robot problems demonstrated the promising performance of the proposed framework.




Fig. 23 Codes for multi-agent reinforcement learning

Impact: The key outcome of this capability is active sensing techniques that go beyond optimization for computational efficiency and information gain that dynamically adjust their perception and motion to achieve resilient situational awareness in the presence of sensor failures, jammed communications, detection risk, and/or compromised agents. As teams of robots are deployed in unknown or partially known failure-prone and adversarial environments, the methods developed here will enable efficient and resilient missions such as ISR and target tracking.

3.3.1.2 Wireless Communication Networks for Distributed Collaborative Systems

The scenarios envisioned and studied in DCIST are of a distributed and collaborative nature. In any distributed and collaborative scenario, autonomous agents must share information among themselves in order to accomplish a common task. The information exchanged can range from sensing information, to agent state

information or even abstract data related to the on-board algorithms being executed by the agents. Whichever case it might be, the agents must ultimately rely on sharing information, a procedure that they necessarily must do over a wireless channel. To this end, this task deals with the design of scalable and adaptive networking algorithms in order to support the DCIST scenarios.

The overall objective of this capability is to design algorithms capable of supporting the required wireless network connectivity necessary for accomplishing the distributed tasks put forward in DCIST. We must begin by noting that existing WiFi and 5G capabilities are not well tailored to DCIST environments. Both of these technologies rely heavily on the availability of infrastructure, which is not a realistic assumption in DCIST environments. Rather, our DCIST agents must operate in environments where communication infrastructure is not available. To bridge this gap, we advocate a three-pronged approach in which we 1) dedicate a part of the team to establish communication infrastructure on demand. We do that by leveraging mobility. 2) Operate in environments with intermittent communications. We do that by leveraging communication opportunities. 3) Optimize team trajectories to acquire information that can be successfully relayed back to the command center. We do that by finding navigation policies based on information gains.

We developed a platform named Mobile Wireless Infrastructure on Demand (Mox 2020). This consists of a framework capable of providing wireless connectivity to multi-robot teams via autonomously reconfiguring ad-hoc networks. In many cases, previous multi-agent systems either assumed the availability of existing communication infrastructure or were required to create a network in addition to completing their objective. Instead, this system explicitly assumes the responsibility of creating and sustaining a wireless network capable of satisfying end-to-end communication requirements of a team of agents, called the task team, performing an arbitrary objective. To accomplish this goal, we use a joint optimization framework that alternates between finding optimal network routes to support data flows between the task agents and improving the performance of the network by repositioning a collection of mobile relay nodes referred to as the network team. In order to verify the operation of the Mobile Wireless Infrastructure on Demand system, we implemented the system on a custom-built quadrotor platform, equipped with conventional IEEE 802.11 WiFi and tested their performance in a large-scale experimental setup as shown in Fig. 24. In this considered environment, the task agents perform a circular patrol of a diameter of around 30 m. This range is sufficient for direct communication between the task agents at the required rate to be impossible. Throughput and delay measurements of this scenario are shown in the figure. The agents progressively move away to

their patrol radius, losing connectivity at around 75 s into the experiment, the moment in which the mobile infrastructure team is activated at around. As the network team comes into play, the throughput and delay are stabilized to a reliable rate (Mox 2020).



Fig. 24 Large-scale experimental setup to test our custom infrastructure

Despite our best efforts to develop infrastructure, we still expect that DCIST teams may need to operate under intermittent communication. We have therefore advanced the design of time-varying networks where the communication links between nodes may emerge and disappear over time (Yu 2020). The objective is to develop motion control and coordination strategies for robots in a team to maintain an intermittently connected communication network while ensuring successful propagation of information throughout the network. In this work, we leverage a robot’s mobility to expand the amount of space a team can monitor while maintaining the connectivity of the mobile robot communication network. The idea is to design coordination strategies that allow robots to “lock in” future periodic encounters with other robots in the team at predetermined locations to enable information exchange (Fig. 25). Outside of these prescheduled encounters, robots can move outside of each other’s communication ranges, enabling the team to achieve wider coverage of a region than if they were required to maintain a fully connected network all the time. Our work shows that mobile communication networks with time-varying connectivity can be formed by synchronizing the frequency in which pairs of robots periodically encountered one another in the workspace. By ensuring future encounter events between pairs of robots, the team not only achieves a communication network with intermittent connectivity, but information can also propagate from any robot to the whole team through the network within a finite time window. Additionally, we have extended this strategy to enable robot teams to form time-varying communication networks that are resilient to the presence of malicious agents (i.e., agents that provide incorrect and/or faulty information). In this context, robots can further leverage their mobility to change the resilience of their networks by actively seeking out new encounters and/or removing previously set encounters with other team members (Yu 2020).




Fig. 25 Designing coordination strategies to allow robots to periodically encounter at predetermined location

A fundamental aspect of our work on mobile infrastructure on demand and intermittent communications is the design of trajectories that improve communication between members of the team. We have also addressed the complementary problem of designing optimal exploration trajectories by a robot team that must maintain communication (Schack 2021). The robot team maximizes the information gain along a path while maintaining communication with other team members and a static base station. Walls or other occluding objects may limit radio communication. Furthermore, the information gain at any one configuration depends on previous positions visited by the robots (i.e., observations at different positions are not independent), so metrics for exploration (e.g., information gain) are non-Markovian in terms of robots' positions. Our approach finds a locally optimal solution in the three steps shown in the scheme in Fig. 26. First, we find a heuristic final robot configuration where the robots can communicate and with a non-zero amount of information gain. Second, we use bidirectional, sampling-based planning to find a satisficing path to this final configuration. Third, we locally optimize the satisficing path. By considering the entire path during optimization, we improve information gain relative to path cost compared to other approaches that consider information gain at only the final configuration. In our tested scenarios, our approach achieved 2–5 times more information gain relative to path cost than baseline sampling-based approaches.




Fig. 26 Developing a locally optimal solution in three steps

Impact: The key objective of this capability is to design a system that dynamically adjusts to achieve resilient communication and co-optimized network and mobility at scale. It also explores questions such as can increased resilience be achieved at the cost of intermittent connectivity? As a team of heterogeneous agents is deployed on an ISR task, they must necessarily share information. The information can range from sensing information, to agent state information, or even abstract data related to the on-board algorithms being executed by the agents. Methods here will determine the configuration of a communication team that can best support a task team, learn optimal resource allocations, and develop vehicle motion coordination strategies to synthesize resilient mobile robot communication networks with time-varying connectivity.

3.3.1.3 Heterogeneity for Resilience

This task complements the Adaptive Swarm Behaviors for Uncertainty Mitigation capability by developing a theory of resilient networking and cooperation that ensures mission progress and preserves core capabilities of the team in the face of failures of whole subnetworks of agents due to disruption in sensing and control, sporadic or permanent communication loss, GPS outage, or degrading visual conditions. The main innovation is the development of mathematical models of resilience, with a particular focus on heterogeneity and scale. The task develops abstractions for information flow, network reconfiguration, and robustness to provide quality of service guarantees for the team's mission by exploiting the heterogeneity in team composition.

We have formulated the problem of resilient multi-robot target tracking with adversarial sensing and communication attacks (Fig. 27). We consider the robots may encounter any fixed number of worst-case sensing and communication attacks from an adversary. The sensing attack on a robot results in the removal of all its sensor measurements to the targets it can observe. The communication attack cuts off the communication links among robots and thus disables the sharing of sensor measurements. Our objective is to investigate resilient, active target tracking algorithms to enable provably good tracking performance, measured by the uncertainty in targets' positions, despite the sensing and communication attacks from the adversary. To this end, we model this resilient target tracking problem as a Stackelberg game or a two-stage sequential game with perfect information between the robots and the adversary. Specifically, the robots play as the leader and plan motions to optimize the tracking performance. While the attacker, as the follower, responds with the worst-case sensing and communication attacks to undermine the tracking performance. With a view to finding the equilibrium of the game, we design a resilient approximation algorithm that provides hard guarantees to secure the team's tracking performance even though the adversary blocks some robots' sensing measurements and blanks some communication links among them.




Fig. 27 A team of robots (green unmanned aerial vehicles [UAVs]) is tasked to track multiple targets (red UAVs). The adversary can attack the sensing/observations (black arrows) of the robots and blank communications (blue dotted lines) between them. The objective is to minimize the uncertainty (light blue ellipses) of the robots' positions despite the sensing and communication attacks.

We have also developed a control framework (Mayya 2021) that implicitly addresses the competing objectives of performance maximization and sensor preservation (which impacts the future performance of the team). Our framework (Fig. 28) consists of a predictive component, which accounts for the anticipated risk, and a reactive component, which maximizes the performance of the team regardless of the failures that have already occurred. We apply this in a scenario where a team of robots with heterogeneous sensors must track a set of hostile targets that induce sensory failures on the robots. The likelihood of failures depends on the

proximity between the targets and the robots. Based on a measure of the abundance of sensors in the team, our framework can generate aggressive and risk-averse robot configurations to track the targets. Crucially, the heterogeneous sensing capabilities of the robots are explicitly considered in each step, allowing for a more expressive risk-performance trade-off.




Fig. 28 Hostile target tracking using a proximity-based model

We have developed a framework for resilience in a networked heterogeneous multi-robot team subject to resource failures (Ramchandaran 2021). Each robot in the team is equipped with resources that it shares with its neighbors. Additionally, each robot in the team executes a task, whose performance depends on the resources to which it has access. When a resource on a particular robot becomes unavailable (e.g., a camera ceases to function), the team optimally reconfigures its communication network so that the robots affected by the failure can continue their tasks. We focus on a monitoring task, where robots individually estimate the state of an exogenous process. We encode the end-to-end effect of a robot’s resource loss on the monitoring performance of the team by defining a new stronger notion of one-hop observability. By abstracting the impact that low-level individual resources have on the task performance, our framework leads to the principled reconfiguration of information flow in the team to effectively replace the lost resource on one robot with information from another. A controller based on finite-time convergence control barrier functions drives each robot to a spatial location that enables the communication links of the modified graph. We validate the effectiveness of our framework (Fig. 29) by deploying it on a team of differential-drive robots estimating the position of a group of quadrotors.




Fig. 29 Experiments with a resilient networked testbed: A team of five quadrotors executing coordinated motion are monitored by ground robots

Impact: The key goal of this capability is to develop a theory of networking and cooperation, with a focus on heterogeneity and scale that ensures mission progress and preserves core capabilities of the team in the face of failures of whole subnetworks of agents due to disruption in sensing and control, communication loss, GPS outage, or degrading visual conditions. As robot teams are deployed to perform a complex breach operation, this project explores the potential trade-offs in a heterogeneous team approach that can distribute required capabilities (sensing, comms, mobility, etc.) amongst the team in order to adapt to sudden and large mission, system, and environmental changes.

3.3.1.4 Adaptive Swarm Behaviors for Uncertainty Mitigation

The objective of this task is to develop tools and strategies that will enable a team of autonomous systems to effectively engage and interact with an opposing team. Such tools are especially necessary for Force Protection and Force Multiplication to enable not only a rich understanding of dynamically evolving and potentially hazardous environments but also the ability to adapt to and engage with increasing uncertainty, infrastructure failures, or adversarial deception by autonomously inferring their intents, taking informative actions to improve the situational awareness of the team, and executing suitable countermeasures. To address these challenges, we have explored the following:

- MARL algorithms that can empower the agents with resilience in dynamic environments or in the presence of adversaries who may rely on deception or other intelligent strategies at mission time. Fast adaptation behaviors necessary to achieve high-performance under such non-stationarity and high sample-efficiency of the MARL algorithm for large team sizes are desirable (Kim 2020; Sun 2020, 2021).

- Distributed strategies for heterogeneous teams to adapt to changing task requirements. The focus is on strategies that can track the evolution of the tasks and determine how best to reallocate or re-task the team’s resources to better achieve the mission objectives (Salam 2022).
- Data-driven system identification for inferring an adversarial swarm’s behavioral patterns and intra-swarm interactions that give rise to the observed behaviors (Zhang 2021).
- Tools to track and engage an adversarial team while managing uncertainty in strategy, execution, and environment for multi-robot surveillance and perimeter security scenarios.

Innovation and Technical Approach. To ensure multi-agent systems are resilient in dynamically changing environments, we leverage latent conditional policy learning with a hierarchical structure, which allows for maximal exploitation of different environmental modes while enabling high adaptability due to the flexible policy network structure. To overcome the non-stationarity from learning and deceptive adversaries, we develop approaches to perceive the intelligent adversaries’ behavior, followed by meta-learning algorithms in the category of fine-tuning methods and beyond. To improve the scalability of such multi-agent learning algorithms with large team size, efficient communication networks and diverse skills are learned to accomplish high sample-efficiency and performance. To enable the adaptation to changing task requirements, we leverage the spectral properties of kernel transfer operators to develop efficient feature-based representations of time-varying density functions in complex environments. The proposed framework allows for data-driven distributed estimation, tracking, and model representation to enable a heterogeneous team to continuously adapt to changing mission requirements and environmental conditions. Simultaneously, data-driven system identification strategies to infer a swarm’s behavior and intent enable effective interactions and responses to an unknown, potentially adversarial, swarm. Techniques to synthesize countermeasures for interacting with adversarial entities in the context of multi-robot surveillance and perimeter defense are considered. The developed strategies focus on approaches that can simultaneously manage uncertainty in strategy, execution, and environment.

We developed a framework to enable a team of heterogeneous mobile robots to model and sense a multiscale system, for example, an animal or human swarm moving through a complex environment (Fig. 30). We propose a coupled strategy, where robots of one type collect high-fidelity measurements at a slow time scale and robots of another type collect low-fidelity measurements at a fast time scale, for the purpose of fusing measurements together. The multiscale measurements are

fused to create a model of a complex, nonlinear process that is dynamically changing across both space and time. The model helps determine optimal sensing locations and predict the evolution of the process. The key contributions are 1) consolidation of multiple types of data into one cohesive model, 2) fast determination of optimal sensing locations for mobile robots, and 3) adaptation of models online for various monitoring scenarios (Salam 2022). We have illustrated the proposed framework by modeling and predicting the evolution of a simulated spatiotemporal process and propose to extend the strategy to track the movement of swarms of agents.





Fig. 30 Example of a heterogeneous team tracking the movement of a herd of animals

In a second sub-task we consider the task of learning swarm behaviors based on the observation of the individual agents' trajectories. Collective swarming behaviors are the results of agent-level dynamics. Extracting these agent-level dynamics is of paramount importance to understanding the emergence of swarming patterns, informing artificial swarm design, and staging adversarial attacks on swarms. However, more often only the observation of swarming trajectories is available, posing a challenge to identifying the agent-level dynamics. We adopt a state-of-the-art continuous-time modeling approach, the knowledge-based neural ordinary differential equations (K-NODEs), to extract agent-level dynamics from observations. The continuous-time nature of K-NODEs enables straightforward knowledge embedding into neural networks for hybrid learning, which drastically reduces the amount of data needed for training and improves the model performance. Using flocking as an example, we apply K-NODEs on a small swarm of flocking agents, and incorporate simple assumptions as knowledge. Our assumptions include a decentralized information structure, a dynamic communication network, and swarm homogeneity, all of which are reasonable assumptions in both natural and artificial swarms. We have demonstrated efficient and scalable learning of the closed-loop agent-level dynamics with K-NODEs,

which finishes training within minutes on the CPU of a desktop computer. The resulting agent-level dynamics model was applied to a larger swarm and the same flocking pattern emerges, demonstrating the generalizability of the learnt agent-level model as shown in Fig. 31. In addition, if the open-loop dynamics for each agent are known, it can be incorporated into the model. In other words, learning decentralized controllers in a swarm is a special case of our learning problem, where the open-loop dynamics can simply be treated as knowledge. To our knowledge, this work is the first adopter of the most recent continuous-time deep learning techniques in the system identification of swarming behaviors (Zhang 2021).



Applying the trained closed-loop dynamics on a swarm with (a) the same number of agents ($N=10$) as the training data, and (b) more agents ($N=50$) than the training data. It can be seen that in both cases, flocking behavior emerges in the swarm. For the swarm with 50 agents, sub-swarms form and flock. This is similar to if we simulate such a large swarm with the ground truth flocking controller.

Fig. 31 Learning to swarm using K-NODEs

In many real-world scenarios, it is imperative for mobile robots to safely navigate through rough terrain. While incorporating geographic and geometric terrain features for these tasks are very important, these aspects have rarely been studied in the multi-robot setting. We have developed a strategy to characterize terrains based on their complexity, and visibility constraints, using metrics adopted from Geographic Information Systems (GIS) and computational geometry respectively. We further developed an approach to show how this characterization can be utilized in multi-robot adversarial settings to formulate optimal strategies for safe and stealthy navigation. This is done in light of a two-player zero-sum game formulation between a heterogeneous team of mobile robots, the transporters, and another stationary team positioned strategically across different locations in the environment, the observers. We have evaluated our strategies using synthetic and real-world terrains (Fig. 32).



Fig. 32 Example of different paths with different detection likelihoods taking into consideration terrain geometry

3.3.1.5 Robust Adaptive Machine Learning

Autonomous learning-enabled teams of robots will need to respond intelligently to unforeseen circumstances. This includes adapting rapidly to changing conditions. By developing robust adaptation techniques that can modify learned models on-the-fly in response to changing mission parameters, environment conditions, and system integrity (e.g., damage to components, sensors, actuators), the algorithms developed as part of this effort will provide a degree of resilience to DCIST platforms that would be difficult to attain with standard static learned models.

To this end, we developed algorithms that enable online adaptation of both individual robot controllers and team-level coordination mechanisms suitable for DCIST platforms, which would cover both perception and control mechanisms. In contrast to prior work in meta-learning and online learning, the focus in this task will specifically be on **robust** and **resilient** adaptation methods that are suitable for multi-agent coordination problems in contested open-world environments. For instance, if one of the robots in a team is immobilized, while the camera on another (still mobile) robot is disabled, the immobilized robot should be able to provide “spotting” capability to the mobile robot with degraded perception. In the domain of meta-learning, a standard assumption in the literature is that the distribution over tasks (i.e., disturbances, commands, etc.) at test-time matches the distribution over which meta-training was performed, so while classically meta-trained models may be able to adapt to new tasks, those tasks themselves must be in-distribution. For practical robust and resilient open-world behavior, this standard assumption is extremely limiting: it is precisely those unpredictable task changes that most require online adaptation. For example, if the meta-training process simulated a variety of

component failure scenarios for a team of robots (e.g., the steering system fails on one robot, a sensor fails on another, etc.), an adversary that understands the adaptive capabilities of the team might intentionally attempt to disable parts of the robot team in such a way as to cause their adaptation mechanism to fail. On-the-fly adaptation to unexpected disturbances is therefore essential. Therefore, we aimed to develop mathematical and algorithmic frameworks for reasoning about adaptation to **unexpected** and **out-of-distribution** tasks, both in a meta-learning setting and in an online learning setting. The first task will focus on the meta-learning setting, and the second task will focus on the online learning setting. Research under this task has developed several key algorithmic advances in terms of basic meta-learning algorithms, and has evaluated these methods in the context of several real-world robotic platforms that reflect capabilities relevant to DCIST. We highlight a few of these algorithms and evaluations here.

Model-based reinforcement learning algorithms learn to perform complex tasks by first learning a predictive model (for example, a neural network that predicts the resulting state when the robot takes a particular action, such as the velocity that will result from applying a particular motor command), and then using this predictive model to construct a plan to solve a given task. In effect, the predictive model represents the algorithm’s estimate of the laws of physics that govern the robot, other objects in the environment, and other agents. However, such algorithms can fail catastrophically at runtime if the learned model fails to generalize to the test-time setting. This could occur, for example, because the robot is operating in a new environment that differs too much from the conditions under which the model was learned, due to mechanical damage to the robot itself, or due to systematic changes in the behavior of other agents. Research under this task developed the first algorithm for model-based meta-reinforcement learning (McAllister 2021), which meta-trains a model that can quickly adapt at test-time to changing environment conditions, damage or changes to the robot itself, and other unexpected events. This approach is fully general, and represents a fundamental advance in model-based reinforcement learning. In experiments, it enables both simulated and real-world robots to adapt to changing dynamics, such as mechanical damage or unexpected terrain, in under a second, whereas naive adaptation methods either fail completely or require several orders of magnitude more data to adapt, precluding adaptation at real-time speeds. The practical implications of such an approach include substantially improved robustness to changing environmental conditions and robots that can dynamically adjust their behavior to handle mechanical damage. In quantitative evaluations, two variants of our method attained final performance that was on some tasks up to $2\times$ better than prior non-meta-learned model-based algorithms and model-free approaches. These results are summarized in Fig. 33.

In future work, analogous methods could also enable autonomous systems that can adapt to changing behavior of other collaborative and adversarial agents.




Fig. 33 Quantitative comparisons of two variants of model-based meta-RL (McAllister 2021) developed as part of this project (gradient-based adaptive learner [GrBAL] and recurrence-based adaptive learner [ReBAL] + model predictive path integral control [MPPI]) on simulated (left) and real-world (middle) benchmark tasks that require rapid adaptation to test-time changes, including mechanical damage to a robot and changing terrains. Right image shows the real-world legged platform used for experiments. The y-axis shows (normalized) total reward over the course of an episode (which includes adaptation time), after meta-trained for a fixed number of samples (chosen to be realistically low, equivalent to a few hours in simulated tasks and 30 min for the real-world tasks). Our methods (GrBAL and ReBAL) attain results that are significantly better than prior approaches, often by $2\times$ or more.

In contrast to model-based reinforcement learning, model-free algorithms learn through trial and error. Such methods can be preferable in settings where learning the “laws of physics” can be difficult—this is especially relevant in multi-agent settings, where learning a model requires also learning how other agents will behave, whereas model-free learning only requires acquiring a suitable strategy. Meta-learning can enhance the efficiency and capability of model-free algorithms as well. Such algorithms are conventionally highly inefficient (for example, the well-known AlphaGo result showed that model-free RL could beat the world champion at Go, but required playing billions of virtual games to learn to do so—this would never be feasible in the real world). Meta-reinforcement learning can in principle learn how to learn via reinforcement, using multiple prior tasks to acquire effective exploration and learning strategies. As part of this research, we developed PEARL (Rakelly 2020), a state-of-the-art model-free meta-reinforcement learning algorithm that improved over the sample efficiency of prior meta-reinforcement learning methods by one to two orders of magnitude depending on the task, as compared to prior algorithms that existed at the time of publication (in 2019). A

graph showing quantitative comparisons to state-of-the-art model-free meta-RL methods at the time of publication on standard benchmark tasks is shown in Fig. 34, with the x-axis corresponding to a log scale in terms of the number of meta-training samples, and the y-axis indicating post-adaptation performance on new tasks at that point in meta-training (which is not available to the algorithm but shown only for evaluation).




Fig. 34 Quantitative comparison of our method, PEARL, with state-of-the-art meta-RL algorithms circa 2019 at the time our method was developed. The x-axis is a log scale showing the number of samples needed by the algorithm, and the y-axis is the performance (total reward) at that point during meta-training on unseen test tasks. Across most benchmarks, PEARL learns about two orders of magnitude faster, and often attains significantly better final results.

Impact: The key outcome of this capability are adaptive machine learning algorithms that enable teams of robots to adapt on-the-fly to unforeseen large and rapid changes in environmental conditions, mission parameters, and robot state by means of online learning and meta-learning algorithms. For example, as a robot or robot team perform complex maneuvers, the agents would use the methods developed here to 1) overcome damage to individual platforms, which may require other robots in the team to realign their objectives to compensate for the diminished capability of one or more of the robots; 2) adapt to novel commands issued by a human teammate or leader; or 3) adapt to changing or unexpected behavior on the part of bystanders or potential adversaries, so as to better predict future behavior and enable a coordinated response.

3.4 Thrust 4: Cross Disciplinary Experiments (CDE)

In the Army's future concept for MDO, robotics and autonomous systems are a key enabler for sensing, detecting, and communicating while operating in contested environments. This task establishes a foundation for operationalizing the use of mixed air-ground robot teams to provide situational awareness at standoff by tying together work in distributed mapping, learning, and coordination to provide an experimental baseline. By proposing metrics that capture some of the essential features of MDO, such as speed, accuracy, and intra-team information latency, and beginning to incorporate adversarial models, such as communications jammers, this task will provide a foundation for future analysis and development programs of the Army.

3.4.1 Capabilities Description

3.4.1.1 CDE-A– Heterogeneous Multi-Agent Situational Awareness

This CDE focuses on a single operational capability of Route Reconnaissance, but evaluates a range of technical capabilities including world modelling, planning, and communication strategies. This specific operation will be used to clarify the requirements for and evaluate the success of a DCIST collaborative framework that will establish 1) easy entry to Army-relevant simulations and experiments; 2) scalable simulation and experimentation in terms of numbers of vehicles, types of vehicles, environment scale and complexity, or level of fidelity/abstraction; 3) access to facilities; and 4) exchange of reference implementations between RA tasks to compose DCIST collaborative capabilities. Toward this end we have already begun implementation of an open testbed to facilitate and drive collaboration among the research teams and to integrate research efforts across the alliance, including our DEVCOM ARL collaborators.

CDE-A explores complex scene understanding, inference and multi-agent planning over long time and length scales in the presence of adversaries in a dynamically changing environment. Particular emphasis is placed on 1) inference of a hybrid semantic-geometric world model; 2) complex mission planning over a distributed team of heterogeneous agents; and 3) resilience in the face of a complex communication environment. These topics will be addressed in the context of the mission profile of Route Reconnaissance.

We consider a scenario where a team is to move through a contested urban area, with multiple routes through the environment and dynamic agents to be avoided. The goal of the blue forces is to detect and identify all static and dynamic obstacles and threats in the scene, while minimizing the detection of the blue forces by the red forces and avoiding regions of conflicts including no-fly zones.

Consider the scenario in Fig. 35, where the team of blue force robots has to move through the environment to the goal location, with stale or possibly no prior information about the environment. The green arrows show some of the example routes that the robots could consider. The robots must collaboratively decide how to partition the problem of exploration for situational awareness to identify routes, as well as locations of the red forces. Scores will be a function of speed of maneuver, as well as the number of red forces detected and the number of blue forces identified.



Fig. 35 A team is moving through a contested urban area, with multiple routes through the environment and where dynamic agents are to be avoided

Individual robots only have local information without line of sight, and the communication network will be intermittent due to environment complexity as well as adversarial jamming. It will be essential to have mobile robots that can operate with only partially situational awareness, and can also merge information when the network is available into a consistent world model.

The scenario can be made incrementally challenging in various axes by varying the area, placing red force agents above the road or inside buildings, increasing the speed and maneuverability of the red force agents, and increasing the complexity of the environment and the availability of communication resources by introducing red force jammers.

Experimental Variables

- 1) Scale: size of area, duration of one episode.
- 2) Complexity of environment: numbers of chokepoints, numbers of types of environments (indoor vs. outdoor, one-story vs. multi-story), the presence of adversarial agents.
- 3) Complexity of communication environment: maximum range, radius of jamming, bandwidth.

- 4) Operational tempo: speed of other agents and changes in the environment
In addition to the four experimental variables, the size and heterogeneity of the team is important.
- 5) Team size: number of robots.
- 6) Heterogeneity: number of types of robots, ground vehicles only, mix of air and ground vehicles.

Evaluation Metrics

- 1) *Time to traverse*: the overall time to complete the site traversal (from start to goal) while obeying constraints (e.g., no-fly zones, staying out of line-of-sight from adversaries). Performance of a human-only team can be used as a baseline for comparison.
- 2) *Latency of awareness*: for a particular piece of information (such as part of an environment map or a detection of an object of interest), the largest length of time between initial acquisition and incorporation into the models of other systems that need it. This captures the effectiveness of distributed intelligence and the ability to incorporate communications and control, and lower numbers are better.
- 3) *Latency of data exfiltration*: a specialization of *latency of awareness*, this refers specifically to the time that it takes for each piece of information to be relayed back to the blue teammates waiting at the starting point.
- 4) *Accuracy of obstacle/threat detection*: the number, correct classification, and accurate estimation of static and dynamic obstacles and threats in the environment. This could apply to all possible obstacles/threats or only to the subset that would impact the future route that the system is seeking to establish.
- 5) *Time-to-detection for objects*: essentially a measure of exploration efficiency, this is the total elapsed time from experiment start until confident detection of objects. This could be normalized against the earliest possible time that an object *could* be detected based on the top speed of the platforms, and lower numbers are better.
- 6) *Time-to-reveal for blue team*: the length of time from experiment start until blue team members are detected by adversaries. Perfectly satisfying the line-of-sight constraints would maximize this time, but generally the goal is to delay this as long as possible.

We have focused on integration of several of the technologies developed elsewhere in the program, notably in Thrust 1 and 2 to validate the performance of these technologies, both in realistic simulation and on field trials. Our Army-relevant convoy protection scenario features a team of air and ground robots traversing an environment containing potential adversaries, where the air robots can be tasked to investigate potential adversaries and determine safe passage for the ground vehicles. The technologies being assessed include the following:

- **OrcVIO:** OrcVIO is the single-platform object-based SLAM technology. The software for this technology has been transitioned into the CDE-A simulator, and work is progressing to optimize it for operation on the real robot, especially in terms of computational speed.
- **Kimera-Multi:** Kimera-Multi is the rich multi-robot metric-semantic mapping system. The software for this technology has been transitioned into the CDE-A simulator and also evaluated on the real CDE-A robot platforms. In the simulator, in the Camp Lejeune environment, we have demonstrated the ability of three robots to build a consistent map simultaneously at $8\times$ speed. We have also demonstrated the ability to recover labeled meshes of buildings (assuming ground-truth semantic labeling).
- **CLEAR and CLIPPER:** CLEAR is the centralized system for alignment of point clouds, and CLIPPER is the robust data association system. Both have been transitioned into the CDE-A simulator. CLEAR has been shown to perform map merging between four robots in simulation, and CLIPPER has been shown to perform inter-robot loop closure detection between air and ground vehicles in simulation. CLIPPER has also been transitioned to real robot platforms, and used to infer coordinate frame alignments between air and ground vehicles.
- **Semantic planning:** The semantic planning system is the technology for using semantic non-geometric information to inform the planner in terms of desirable and undesirable trajectories. The software for this technology has been transitioned into the CDE-A simulator, and trained to infer trajectories that minimize visibility profiles, using the semantics of the environment to predict observer locations and concurrent visibility. Figure 36 gives example trajectories in simulation. This technology has also been transitioned to the real CDE-A quadrotor platform.
- CDE-A has a strong commitment to validating its algorithms in perception, estimation, and planning in photo-realistic simulation. This Unity-based simulation was originally developed at MIT, transitioned to DEVCOM

ARL, and allows several of the tasks under CDE-A to validate performance before assessing on real vehicles.

- CDE-A has begun assessment on real vehicles. Initial trials of mixed quadrotor-ground vehicle systems were conducted in late 2020 at test facilities in both the Boston area and the Philadelphia area. Further extensions of these trials to experimentally assess Kimera-Multi, CLIPPER, and the semantic planner were conducted in May 2021. Figure 37 gives an example mission in this experimental exercise.




Fig. 36 Two example trajectories of a quadrotor choosing minimum visibility trajectories that it has learned to predict from the semantics in the environment




Fig. 37 The test facility at Medfield, Massachusetts, outside Boston. The blue trajectory is the initial trajectory of a Clearpath Jackal, and the green trajectory is the initial trajectory of a quadrotor using the stealthy semantic navigation strategy. When objects of interests are detected by the Jackal, the quadrotor is tasked with deviations to inspect and refine the pose of the objects, denoted by the dashed green lines. This scenario exercises the state estimation, object localization, coordinate frame alignment, and stealth navigation.

3.4.1.2 CDE B – Dynamic Teaming Operations in Contested Environment

To realize the vision of teams of distributed intelligent heterogeneous systems as force multipliers and force protectors for Soldiers in complex military relevant environments, the innovations proposed in each of the research thrust areas must be informed by the needs of large-scale, high op-tempo operations in the presence of adversaries and evaluated through simulation and experimentation in representative environments.

This CDE addresses operational capabilities for teams in adversarial environments with a focus on *Perimeter Defense*. The goal is to defend against a red team of sophisticated adversaries by deploying large numbers of heterogeneous blue team assets in a complex environment with the constrained communications inherent to Army-relevant scenarios. The key innovations include 1) synthesis of team behaviors with high operational tempo to respond and adapt to dynamic adversaries; and 2) establishing the communication networks required for coordination and collaboration among team members.

We developed algorithms for a team of ground robots tasked with intercepting intruders and defending a perimeter (a pre-specified high security zone), and demonstrated the ability to perform swarm-versus-swarm maneuvers using game theoretic techniques in the DCIST simulator. To overcome the limited visibility of these robots operating in occluded environments such as urban canyons, aerial robots equipped with sensors to detect intruders are deployed around the perimeters, with the aim of relaying intruder position information to the defense team. These robots utilize results from the resilient multi-robot coverage research performed in Thrust 3. Figure 38 (right) shows the six defender Warthog robots (shown in blue) instantiated on two high-security perimeters, with aerial intruder robots (shown in red) making their way to the perimeters in the urban canyon.



Fig. 38 (Left) Patrolling operations of 15 Warthog robots in the DCIST Unity simulation. (Right) Perimeter defense operations on two urban perimeters using six Warthog defender agents (blue) and six aerial intruders (red).

We have also developed a Mobile Infrastructure on Demand (MID), a team of agents whose sole task is to provide and guarantee network connectivity across task agents (defender Warthog robots), performing a supporting role critical to successful task accomplishment.

Figure 39 (left) shows a configuration of aerial robots that are dedicated to ensuring communication among the defense team. These agents leverage a connectivity maintenance algorithm to constantly move in order to maintain connectivity among the task agents. The image on the right shows the view of a building from a quadrotor.



Fig. 39 (Left) shows a configuration of aerial robots that are dedicated to ensuring communication among the defense team. These agents leverage a connectivity maintenance algorithm to constantly move in order to maintain connectivity among the task agents. (Right) shows the view of a building from a quadrotor.

While the above-described behaviors aim to intercept intruders in the immediate vicinity of the perimeters, it is also critical to maintain situational awareness over the large operational environment. Since maintaining complete coverage is impossible due to the size of the environment and limitations on the team size, patrolling behaviors are being generated by solving a polymatrix game between the aerial and ground patrol agents. We aim to test the operations of this integrated system by simulating sharp changes in intruder deployments and testing the ability of the robot team to effectively respond to the adversary. Sophisticated mechanisms for simultaneous task allocation, task planning, scheduling, and motion planning will be leveraged to coordinate the operations of the robot team.

4. Key Program Metrics and Impact (18Q3 to 20Q3)

4.1 Education: Number of Students Supported

Undergraduates – 6

Masters – 36

PhDs – 107

Post Docs – 32

4.2 Number of Publications

Journals – 29

Conferences – 237

4.3 Notable Papers, Awards, and Recognitions

- 1) 2020 IEEE ICRA **Best Paper Award** in Robot Vision: “*Graduated Non-Convexity for Robust Spatial Perception: From Non-Minimal Solvers to Global Outlier Rejection*” Heng Yang, Pasquale Antonante, Vasileios Tzoumasand, and Luca Carlone.
- 2) 2020 IEEE ICRA **Finalist for Best Paper Award** in Robot Vision: “Metrically-Scaled Monocular SLAM using Learned Scale Factors” W. Nicholas Greene and Nicholas Roy.
- 3) 2020 ICASSP **Best Paper Award**: "Better Safe than Sorry: Risk-aware Nonlinear Bayesian Estimation" by Dionysios Kalogerias, Luiz Chamon, George J. Pappas and Alejandro Ribeiro.
- 4) 2020 ICASSP **Best Student Paper Award**: "The Empirical Duality Gap of Constrained Statistical Learning," by Luiz Chamon, Santiago Paternain, Miguel Calvo-Fullana, and Alejandro Ribeiro.
- 5) 2019 ICRA **Finalist Best Paper Award** on HRI: “Deconfliction of Motion Paths with Traffic Inspired Rules in Robot–Robot and Human–Robot Interactions,” Federico Celi, Li Wang, Lucia Pallottino, and Magnus Egerstedt.
- 6) 2019 RSS **Best Student Paper**: “An Online Learning Approach to Model Predictive Control,” Wagener, Nolan; Cheng, Ching-an; Sacks, Jacob; Boots, Byron (DCIST Supplemental Task).

- 7) 2019 Eusipco–**Student Paper Award**: “Gated Graph Convolutional Recurrent Neural Networks,” Luana Ruiz, Fernando Gama, and Alejandro Ribeiro.
- 8) 2019 ACC **Finalist Best Student Paper Award**: “Motion Planning with Secrecy,” Anastasios Tsiamis(Student Author), Andreea Alexandru, George J. Pappas.
- 9) 2019 ICCPS **Finalist Best Paper Award**: “Encrypted LQG using Labeled Homomorphic Encryption,” Andreea Alexandru and George Pappas.
- 10) 2019 IROS **Finalist Best Paper Award**: “Safety, Security, and Rescue Robotics, FASTER: Fast and Safe Trajectory Planner for Flights in Unknown Environments,” Jesus Tordesillas Torres, Brett Lopez, and Jonathan Patrick How.
- 11) 2019 ACC **O. Hugo Schuck Best Paper Award**: “Permissive Barrier Certificates for Safe Stabilization Using Sum-of-squares” Li Wang, Dongkun Han, and Magnus Egerstedt.

5. Conclusion

Research within the DCIST CRA has significantly advanced the state of the art in multi-agent autonomy for Army applications in the areas of Multi-Agent and Resilient Situational Awareness, Collaborative Learning and Intelligence, Adaptation and Learning in Wireless Autonomous Systems, Hierarchical Abstractions for Planning, Joint Resource Allocation in Perception-Action-Communication Loops, Hierarchical & Distributed Control for Adversarial Operations, Scalable Task Assignment for Heterogeneous Multi-Unit Teams, Tactical Engagement of Heterogeneous Teams in Complex Environments, Human Interaction with Large Heterogeneous Teams, Heterogeneity for Resilience, and Adaptive Swarm Behaviors for Uncertainty Mitigation.

Future collaborative robotic systems will leverage this research to create and align geometric and semantic information to reduce uncertainty and enable better individual and team localization, mapping, and path planning in complex dynamic environments; adapt to previously unseen events, communication link losses, and changes to the environment; perform fast, on-the-fly, replanning at both the local and global scale; strategic deployments of agents in dynamically changing scenarios involving models of adversarial agents and imperfect/delayed communication; produce sufficient or “good enough” solutions, in a computationally feasible, distributed, and adaptive manner; improve techniques for how humans perceive and communicate time-evolving information; perform

active sensing that dynamically adjusts perception and motion to achieve resilient situational awareness in the presence of sensor failures, jammed communications, detection risk, and/or compromised agents; and achieve resilient communication and co-optimized network and mobility at scale.

6. References

- Atanasov N. A unifying view of geometry, semantics, and data association in SLAM. Proceedings of the International Joint Conference on Artificial Intelligence; 2018.
- Blumenkamp J. The emergence of adversarial communication in multi-agent reinforcement learning. Proceedings of the Conference on Robot Learning (CoRL); 2020.
- Cai X, Schlotfeldt B, Khosoussi K, Atanasov N, Pappas GJ, How JP. Non-monotone energy-aware information gathering for heterogeneous robot teams. Proceedings of the IEEE Conference on Robotics and Automation, Xian; 2021 June.
- Chamon L. Counterfactual programming for optimal control. Proceedings of the 2nd Conference on Learning for Dynamics and Control; 2020.
- Chang Y. Kimera-Multi: a system for distributed multi-robot metric-semantic simultaneous localization and mapping. Proceedings of ICRA; 2020.
- Duong T. Autonomous navigation in unknown environments using sparse kernel-based occupancy mapping. Proceedings of the IEEE International Conference on Robotics and Automation; 2020.
- Emam Y. Adaptive task allocation for heterogeneous multi-robot teams with evolving and unknown robot capabilities. Proceedings of ICRA; 2020.
- Fathian K. CLEAR: a consistent lifting, embedding, and alignment rectification algorithm for multi-view data association. Proceedings of IEEE TRO; 2019.
- Feng Q. Localization and mapping using instance-specific mesh models. Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems; 2019.
- Feng Q. Fully convolutional geometric features for category-level object alignment. Proceedings of IROS; 2020.
- Kahn G. BADGR: An autonomous self-supervised learning-based navigation system. Proceedings of RA-L; 2021a.
- Kahn G. LaND: Learning to navigate from disengagements. Proceedings of RA-L; 2021b.

- Kantaros Y, Pappas GJ. Distributed active information acquisition for multi-robot systems. Proceedings of the IEEE Conference on Robotics and Automation, Xian; 2021 June.
- Khan A. Graph policy gradients for large scale unlabeled motion planning with constraints. Proceedings of IROS; 2020.
- Khan A. Graph policy gradients for large scale robot control. Proceedings of CORL; 2019.
- Kim D. A policy gradient algorithm for learning to learn in multiagent reinforcement learning. <https://arxiv.org/abs/2011.00382>; 2020.
- Kolb J, Kishore M, Shaw K, Ravichandar H, Chernova S. Predicting individual human performance in human-robot teaming. Proceedings of the IEEE International Conference on Robot and Human Interactive Communication; 2021.
- Larsson D. Information-theoretic abstractions for planning in agents with computational constraints. IEEE Robotics and Automation Letters (under review); 2021.
- Lajoie P-Y. DOOR-SLAM: distributed, online, and outlier resilient slam for robotic teams. Proceedings of ICRA & RAL; 2020.
- Liu K. Learned sampling distributions for efficient planning in hybrid geometric and object-level representations. Proceedings of ICRA; 2020.
- Lusk P. CLIPPER: A graph-theoretic framework for robust data association. Proceedings of ICRA; 2020.
- Lynn CW, Bassett DS. Compressibility of complex networks. arXiv preprint (arXiv:2011.08994) Proceedings of the National Academy of Sciences; 2020.
- Malencia M. Fair robust assignment using redundancy. Proceedings of R-AL; 2021.
- Milano F. Primal-dual mesh convolutional neural networks. Proceedings of the Conference on Neural Information Processing Systems; 2020.
- Mayya S. Adaptive and risk-aware target tracking with heterogeneous robot teams, In preparation.
- McAllister R. Model-based meta-reinforcement learning for flight with suspended payloads. <https://arxiv.org/abs/2004.11345>; 2021.

- Mox D. Mobile wireless network infrastructure on demand. Proceedings of International Conference on Robotics and Automation (ICRA); 2020.
- Paritosh P. Hypothesis assignment and partial likelihood averaging for cooperative estimation. Proceedings of the IEEE Conference on Decision and Control; 2019.
- Paritosh P. Marginal density averaging for distributed node localization from local edge measurements. Proceedings of the IEEE Conference on Decision and Control; 2020.
- Paulos J. Decentralization of multiagent policies by learning what to communicate. Proceedings of the International Conference on Robotics and Automation; 2019.
- Pierpaoli P. Inferring and learning multi-robot policies by observing an expert. arXiv:1909.07887; 2020.
- Rakelly K. Efficient off-policy meta-reinforcement learning via probabilistic context variables. <https://arxiv.org/abs/1903.08254>; 2020.
- Ramchandaran R. Resilient monitoring in heterogeneous multi-robot systems through network reconfiguration. IEEE Transactions on Robotics and Automation, conditionally accepted to IEEE Transactions on Robotics (T-RO); 2021.
- Ravichandar H, Shaw K, Chernova S. STRATA: a unified framework for task assignments in large teams of heterogeneous robots. J Autonomous Agents Multi-Agent Syst. 2020.
- Rawls J. A theory of justice. The Belknap Press of Harvard University Press; 1971.
- Rosinol A. Incremental visual-inertial 3D mesh generation with structural regularities. Proceedings of ICRA; 2019.
- Rosinol A. Kimera: from SLAM to Spatial Perception with 3D Dynamic Scene Graphs. IJRR; 2021a.
- Rosinol A. Actionable spatial perception with places, objects, and humans. Robotics: Science and Systems, 2020b.
- Rosinol A. Kimera: an open-source library for real-time metric-semantic localization and mapping. Proceedings of ICRA; 2020c.
- Salam T. Heterogeneous robot teams for modeling and prediction of multiscale environmental processes. submitted to Autonomous Robots, under review. arXiv <https://arxiv.org/pdf/2103.10383.pdf>.

- Schack M. Optimizing non-Markovian information gain under physics-based communication constraints. *IEEE Robotics and Automation Letters*; 2021.
- Shah D. ViNG: Learning open-world navigation with visual goals. *Proceedings of ICRA*; 2021a.
- Shah D. RECON: Rapid exploration for open-world navigation with latent goal models. *5th Annual Conference on Robot Learning*; 2021b.
- Shan M. OrcVIO: Object residual constrained visual-inertial odometry. *Proceedings of IROS*; 2020.
- Shannon CJ, Horney C, Jackson KF, How JP. Human-autonomy teaming using flexible human performance models: An initial pilot study. *Advances in Human Factors in Robots and Unmanned Systems*. Springer International Publishing; 2017. p. 211–224.
- Shishika D. Cooperative team strategies for multi-player perimeter-defense games. *IEEE Robot Automat Lett*; 2020a.
- Shishika D. Game theoretic formation design for probabilistic barrier coverage. *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems*; 2020b.
- Shishika D. Partial information target defense game. *Proceedings of International Conference on Robotics and Automation*; 2021.
- Silva A. Unsupervised role discovery using temporal observations of agents. *Proceedings of International Conference on Autonomous Agents and Multi-Agents Systems*; 2019.
- Sun C. Scaling up multiagent reinforcement learning for robotic systems: learn an adaptive sparse communication graph. *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*; 2020.
- Sun C. FISAR: Forward invariant safe reinforcement learning with a deep neural network-based optimize. *Proceedings of ICRA 2021*.
- Suri N, Breedy MR, Marcus KM, Fronteddu R, Cramer E, Morelli A, Campioni L, Provosty M, Enders C, Tortonesi M, Nilsson J. Experimental evaluation of group communications protocols for data dissemination at the tactical edge. *Proceedings of the IEEE International Conference on Military Communications and Information Systems (ICMCIS)*; 2019.
- Tian Y. Search and rescue under the forest canopy using multiple UAS. *Proceedings of the International Symposium on Experimental Robotics*; 2018.

- Tian Y. Distributed certifiably correct pose-graph optimization. *IEEE Trans Robot.* 2020.
- Tolstaya E. Learning decentralized controllers for robot swarms with graph neural networks. *Proceedings of RSS*; 2020.
- Tolstaya E. Multi-robot coverage and exploration using spatial graph neural networks. *Proceedings of IROS*; 2021.
- Tsiotras P. Bounded rationality in learning, perception, decision-making, and stochastic games. *Handbook of Reinforcement Learning and Control*; 2021.
- Wang B, Xie J, Atanasov N. Coding for distributed multi-agent reinforcement learning. <https://arxiv.org/abs/2101.02308>; 2021.
- Yu X. Synthesis of a time-varying communication network by robot teams with information propagation guarantees. *IEEE Robot Automat Lett.* 2020.
- Zhang J. Knowledge-based learning of nonlinear dynamics and chaos. *Forthcoming*. <https://arxiv.org/abs/2010.03415>.
- Zobeidi E. Dense incremental metric-semantic mapping via sparse Gaussian process regression. *Proceedings of IROS*; 2020.

Appendix. DCIST Program Notable Publications (18Q3–21Q3)

DCIST Program to Date Bibliography Feb 26, 2022

1. Title: A Unifying View of Geometry, Semantics, and Data Association in SLAM, Venue: International Joint Conference on Artificial Intelligence (IJCAI), Lead Author: Nikolay Atanasov, Year: 2018
2. Title: Adversarial Information Acquisition, Venue: Robotics: Science and Systems Workshop, Lead Author: B. Schlotfeldt, Year: 2018
3. Title: Aggregation Graph Neural Networks, Venue: ICASSP, Lead Author: Fernando Gama, Year: 2019
4. Title: Attention and Anticipation in Fast Visual-Inertial Navigation, Venue: Transactions on Robotics, Lead Author: Luca Carlone, Year: 2018
5. Title: Collective Online Learning of Gaussian Processes in Massive Multi-Agent Systems, Venue: Thirty-Third AAAI Conference on Artificial Intelligence (AAAI-19), Lead Author: Trong Nghia Hoang, Year: 2019
6. Title: Composable Learning with Sparse Kernel Representations, Venue: International Conference on Intelligent Robots and Systems, Lead Author: Ekaterina Tolstaya, Year: 2018
7. Title: Control Aware Radio Resource Allocation in Low Latency Wireless Control Systems, Venue: IEEE Internet of Things Journal, Lead Author: Mark Eisen, Year: 2019
8. Title: Convolutional Neural Networks via Node-Varying Graph Filters, Venue: 2018 IEEE Data Science Workshop, Lead Author: Fernando Gama, Year: 2018
9. Title: Coordinating Multi-Robot Systems Through Environment Partitioning For Adaptive Informative Sampling, Venue: ICRA, Lead Author: Nikolas Fung, Year: 2019
10. Title: Coverage Control for Multi-Robot Teams with Heterogeneous Sensing Capabilities, Venue: Robotics and Automation Letters, Lead Author: Maria Santos, Year: 2018
11. Title: Coverage Control for Multi-Robot Teams with Heterogeneous Sensing Capabilities Using Limited Communications, Venue: IEEE/RSJ International Conference on Intelligent Robots and Systems, Lead Author: Maria Santos, Year: 2018
12. Title: Deconfliction of Motion Paths with Traffic Inspired Rules in Robot-Robot and Human-Robot Interactions, Venue: RA-L, IEEE Robotics &

- Automation Letters, and ICRA, IEEE International Conference on Robotics and Automation, Lead Author: Federico Celi, Year: 2019
13. Title: Dense Spatial Segmentation from Sparse Semantic Information, Venue: Workshop on Learning and Inference in Robotics at RSS, Lead Author: Qiaojun Feng, Year: 2018
 14. Title: Learning in Non-Stationary Wireless Control Systems via Newton's Method, Venue: American Controls Conference, Lead Author: Mark Eisen, Year: 2018
 15. Title: Learning In Wireless Control Systems Over Non-Stationary Channels, Venue: IEEE Transaction on Signal Processing, Lead Author: Mark Eisen, Year: 2018
 16. Title: Learning Statistically Accurate Resource Allocations In Non-Stationary Wireless Systems, Venue: International Conference on Acoustics, Speech, and Signal Processing, Lead Author: Mark Eisen, Year: 2018
 17. Title: Local-Game Decomposition for Multiplayer Perimeter-Defense Problem, Venue: CDC, Conference on Decision and Control, Lead Author: Daigo Shishika, Year: 2018
 18. Title: Localization, Grasping, and Transportation of Magnetic Objects by a team of MAVs in Challenging Desert-Like Environments, Venue: IEEE Robotics and Automation Letters and ICRA 2018, Lead Author: Giuseppe Loianno, Year: 2018
 19. Title: LQG Control and Sensing Co-design, Venue: Transactions on Automatic Control, Lead Author: Vasileios Tzoumas, Year: 2020
 20. Title: Meta-Learning Through Coupled Optimization in Reproducing Kernel Hilbert Spaces, Venue: ACC, Lead Author: Juan Cerviño, Year: 2019
 21. Title: MIMO Graph Filters for Convolutional Networks, Venue: 2018 19th IEEE International Workshop on Signal Processing for Advances in Wireless Communications, Lead Author: Fernando Gama, Year: 2018
 22. Title: On the Convergence of Distributed Subgradient Methods under Quantization, Venue: 2018 56th Annual Allerton Conference on Communication, Control, and Computing (Allerton), Lead Author: Think Doan, Year: 2018
 23. Title: On the Trade-Off Between Communication and Execution Overhead for Control of Multi-Agent Systems, Venue: ACC, American Control Conference, Lead Author: Anqi Li, Year: 2019

24. Title: Online Deep Learning in Wireless Communication Systems, Venue: Asilomar Conference on Signals, Systems and Computers, Lead Author: Mark Eisen, Year: 2018
25. Title: Optimal Covariance Control for Stochastic Systems Under Chance Constraints, Venue: 57th IEEE Conference on Decision and Control, Lead Author: Kazuhide Okamoto, Year: 2018
26. Title: Optimization of Switched Linear Systems Over Non-Stationary Wireless Channels, Venue: International Workshop on Signal Processing Advances in Wireless Communications, Lead Author: Mark Eisen, Year: 2018
27. Title: Predicting Power Outages Using Graph Neural Networks, Venue: 2018 6th IEEE Global Conference on Signal and Information Processing, Lead Author: Damian Owerko, Year: 2018
28. Title: Prioritized Path Planning in Heterogeneous Robot Teams, Venue: ICRA, Lead Author: Wenying Wu, Year: 2020
29. Title: Probabilistic Model-Agnostic Meta-Learning, Venue: NIPS, Lead Author: Chelsea Finn, Year: 2018
30. Title: Resilient Backbones in Hexagonal Robot Formations, Venue: DARS, International Symposium on Distributed Autonomous Robotic Systems, Lead Author: David Saldana, Year: 2018
31. Title: Resilient Non-Submodular Maximization over Matroid Constraints, Venue: IEEE Transactions on Automatic Control, Lead Author: Vasileios Tzoumas, Year: 2021
32. Title: Resource-Aware Algorithms for Distributed Loop Closure Detection with Provable Performance Guarantees, Venue: International Workshop on the Algorithmic Foundations of Robotics, Lead Author: Yulun Tian, Year: 2018
33. Title: Sample Complexity of Networked Control Systems Over Unknown Channels, Venue: 57th IEEE Conference on Decision and Control, Lead Author: Konstantinos Gatsis, Year: 2018
34. Title: VIO-Swarm: An Autonomous Swarm of Vision Based Quadrotors, Venue: IEEE International Conference on Robotics and Automation ICRA 2018, Workshop Robot Teammates Operating in Dynamic, Unstructured Environments (RT-DUNE), Lead Author: Aaron Weinstein, Year: 2018
35. Title: Visual Inertial Odometry Swarm: An Autonomous Swarm of Vision-Based Quadrotors, Venue: IEEE Robotics and Automation Letters and ICRA 2018, Lead Author: Aaron Weinstein, Year: 2018

36. Title: 3-Dimensional Keypoint Repeatability for Heterogeneous Multi-Robot SLAM, Venue: IEEE International Conference on Robotics and Automation (ICRA), Lead Author: Elizabeth Boroson, Year: 2019
37. Title: Autonomous Landing On a Moving Vehicle with an Unmanned Aerial Vehicle, Venue: Journal of Field Robotics, Lead Author: Tomas Baca, Baca Petr, Stepan Vojtech Spurny, Daniel Hert, Robert Penicka, Martin Saska, Justin Thomas, Giuseppe Loianno, Vijay Kumar, Year: 2019
38. Title: Control Aware Communication Design for Time Sensitive Wireless Systems, Venue: ICASSP, Lead Author: Mark Eisen, Year: 2019
39. Title: Convergence Rates of Distributed Gradient Methods under Random Quantization, Venue: IEEE Trans. Automatic Control, Lead Author: Thinh Doan, Year: 2021
40. Title: Convolutional Neural Network Architectures for Signals Supported on Graphs, Venue: IEEE Transactions on Signal Processing, Lead Author: Fernando Gama, Year: 2019
41. Title: Cooperative Autonomous Search, Grasping, and Delivering in a Treasure Hunt Scenario by a Team of Unmanned Aerial Vehicles, Venue: Journal of Field Robotics, Lead Author: Vojtěch Spurný, Year: 2018
42. Title: Design Guarantees for Resilient Robot Formations on Lattices, Venue: IEEE Robotics and Automation Letters, Lead Author: Luis Guerrero, Year: 2019
43. Title: Diffusion Scattering Transforms on Graphs, Venue: 2019 7th International Conference on Learning Representations, Lead Author: Fernando Gama, Year: 2019
44. Title: Efficient Trajectory Planning for High Speed Flight in Unknown Environments, Venue: ICRA 2019, Lead Author: Markus Ryll, Year: 2019
45. Title: Inertial Velocity and Attitude Estimation for Quadrotors, Venue: IEEE/RSJ International Conference on Intelligent Robots and Systems IROS 2018, Lead Author: James Svacha, Year: 2018
46. Title: Knowledge Gaps in the Early Growth of Semantic Feature Networks, Venue: Nature Human Behavior, Lead Author: Ann Sizemore Blevins, Year: 2018
47. Title: Latency-Reliability Tradeoffs for State Estimation, Venue: TAC, Lead Author: Konstantinos Gatsis, Year: 2020

48. Title: Learning Implicit Sampling Distributions for Motion Planning, Venue: IROS 2018, Lead Author: Clark Zhang, Year: 2018
49. Title: Learning Models of Sequential Decision-Making with Partial Specification of Agent Behavior, Venue: AAAI Conference on Artificial Intelligence, Lead Author: Vaibhav Unhelkar, Year: 2019
50. Title: Locally adaptive kernel estimation using sparse functional programming, Venue: Asilomar 2018, Lead Author: Maria Peifer, Year: 2018
51. Title: Model Predictive Trajectory Tracking and Collision Avoidance for Reliable Outdoor Deployment of Unmanned Aerial Vehicles, Venue: IEEE/RSJ International Conference on Intelligent Robots and Systems IROS 2018, Lead Author: Tomas Baca, Year: 2018
52. Title: Modeling Perceptual Aliasing in SLAM via Discrete-Continuous Graphical Models, Venue: RA-L & ICRA, Lead Author: Pierre-Yves Lajoie, Year: 2019
53. Title: Network Constraints on Learnability of Probabilistic Motor Sequences, Venue: Nature Human Behavior, Lead Author: Ari E. Kahn, Year: 2018
54. Title: Nuclear Environments Inspection with Micro Aerial Vehicles: Algorithms and Experiments, Venue: International Symposium on Experimental Robotics ISER, Lead Author: Dinesh Thakur, Year: 2018
55. Title: Optimal Task Distribution in Heterogeneous Multi-Robot Systems, Venue: European Control Conference, Lead Author: Gennaro Notomista, Year: 2019
56. Title: Overcoming Blind Spots in the RealWorld: Leveraging Complementary Abilities for Joint Execution, Venue: AAAI Conference on Artificial Intelligence, Lead Author: Ramya Ramakrishnan, Year: 2019
57. Title: Resilient Active Information Gathering with Mobile Robots, Venue: IEEE/RSJ International Conference on Intelligent Robots and Systems, Lead Author: B. Schlotfeldt, Year: 2018
58. Title: Resilient Monotone Sequential Maximization, Venue: Proceedings of the 57th IEEE Conference on Decision and Control (CDC), Lead Author: Vasileios Tzoumas, Year: 2018
59. Title: Search and Rescue under the Forest Canopy using Multiple UAS, Venue: International Symposium on Experimental Robotics, Lead Author: Yulun Tian, Year: 2018

60. Title: Self-Assembly of a Class of Infinitesimally Shape-Similar Frameworks, Venue: IEEE/RSJ International Conference on Intelligent Robots and Systems, Lead Author: Ian Buckley, Year: 2019
61. Title: Spatio-Temporally Smooth Local Mapping and State Estimation Inside Generalized Cylinders With Micro Aerial Vehicles, Venue: IEEE Robotics and Automation Letters and IROS 2018, Lead Author: Tolga Ozaslan, Year: 2018
62. Title: Accelerated Inference in Markov Random Fields via Smooth Riemannian Optimization, Venue: RA-L & ICRA, Lead Author: Siyi Hu, Year: 2019
63. Title: Active Perception in Adversarial Scenarios using Maximum Entropy Deep Reinforcement Learning, Venue: ICRA, Lead Author: Macheng Shen, Year: 2019
64. Title: Adaptive Sampling and Reduced Order Modeling of Dynamic Processes by Robot Teams, Venue: IEEE RA-L, Lead Author: Tahiya Salam, Year: 2019
65. Title: Asymptotic Optimality of a Time Optimal Path Parametrization Algorithm, Venue: IEEE Control Systems Letters, Lead Author: Igor Spasojevic, Year: 2019
66. Title: Block-Coordinate Minimization for Large SDPs with Block-Diagonal Constraints, Venue: Technical Report (arXiv), Lead Author: Yulun Tian, Year: 2019
67. Title: Dual Domain Learning of Optimal Resource Allocations in Wireless Communication Systems, Venue: ICASSP 2019, Lead Author: Mark Eisen, Year: 2019
68. Title: Dynamic Tube MPC for Nonlinear Systems, Venue: American Controls Conference, Lead Author: Brett Lopez, Year: 2019
69. Title: Incremental Visual-Inertial 3D Mesh Generation with Structural Regularities, Venue: ICRA, Lead Author: Antoni Rosinol, Year: 2019
70. Title: Large Scale Wireless Power Allocation with Graph Neural Networks, Venue: SPAWC 2019, Lead Author: Mark Eisen, Year: 2019
71. Title: Learning Decentralized Controllers for Robot Swarms with Graph Neural Networks, Venue: International Conference on Robot Learning, Lead Author: Ekaterina Tolstaya, Year: 2020

72. Title: Learning Optimal Resource Allocations in Wireless Systems, Venue: IEEE Transactions on Signal Processing, Lead Author: Mark Eisen, Year: 2019
73. Title: Median Activation Functions for Graph Neural Networks, Venue: ICASSP 2019, Lead Author: Luana Ruiz, Year: 2019
74. Title: Navigation of a Quadratic Potential with Ellipsoidal Obstacles, Venue: CDC, Lead Author: Harshat Kumar, Year: 2019
75. Title: Redundant Robot Assignment on Graphs with Uncertain Edge Costs, Venue: DARS, Lead Author: Amanda Prorok, Year: 2019
76. Title: Robot Co-design: Beyond the Monotone Case, Venue: ICRA, Lead Author: Luca Carlone, Year: 2019
77. Title: S. Kemna, G. Sukhatme, Coordinating Multi-Robot Systems through Environment Partitioning for Adaptive Informative Sampling, ICRA, May 2019, Montreal, Venue: ICRA, Lead Author: Nick Fang, Year: 2019
78. Title: Task Allocation for Heterogeneous Multi-Robot Teams, Venue: IROS, Lead Author: Yousef Emam, Year: 2019
79. Title: 6D Interaction Control with Aerial Robots: The Flying End-Effector Paradigm, Venue: The International Journal of Robotics Research, Lead Author: Markus Ryll, Year: 2019
80. Title: A Decentralized Heterogeneous Control Strategy for a Class of Infinitesimally Shape-Similar Formations, Venue: ICRA, Lead Author: Ian Buckley, Year: 2019
81. Title: A Taxonomy for Characterizing Modes of Interactions in Goal-driven, Human-robot Teams, Venue: IROS, Lead Author: Priyam Parashar, Year: 2019
82. Title: Activity Recognition by Learning from Human and Object Attributes, Venue: International Conference on Robotics and Automation (ICRA) Workshop on Robot Teammates Operating in Dynamic, Unstructured Environments (RT-DUNE), Lead Author: B. Reily, Year: 2019
83. Title: Adaptive Sampling and Energy Efficient Navigation in Time-Varying Flows, Venue: Autonomous Underwater Vehicles: Design and Practice, Lead Author: Tahiya Salam, Year: 2019
84. Title: All Graphs Lead to Rome: Learning Geometric and Cycle-Consistent Representations with Graph Convolutional Networks, Venue: CVPR 2019

- Workshop Image Matching: Local Features and Beyond, Lead Author: Stephen Phillips, Year: 2019
85. Title: An Online Learning Approach to Model Predictive Control, Venue: Proceedings of Robotics Science and Systems XV (RSS), Lead Author: Nolan Wagener, Year: 2019
 86. Title: An Optimal Task Allocation Strategy for Heterogeneous Multi-Robot Systems, Venue: European Control Conference, Lead Author: Gennaro Notomista, Year: 2019
 87. Title: Asymptotically Optimal Planning for Non-myopic Multi-Robot Information Gathering, Venue: Robotics: Science and Systems (RSS), Lead Author: Yiannis Kantaros, Year: 2019
 88. Title: Bayesian-Markov Feedback in Constraint-based Planning, Venue: International Conference on Robotics and Automation (ICRA) Workshop on Robot Teammates Operating in Dynamic, Unstructured Environments (RT-DUNE), Lead Author: M. Schack, Year: 2019
 89. Title: Channels in High-Fidelity Simulations of Unmanned Aerial Systems, Venue: Signal Processing Advances in Wireless Communications (SPAWC), Lead Author: T. R. Godbole, Year: 2019
 90. Title: Convergence Rates of Distributed Two-Time-Scale Gradient Methods Under Random Quantization, Venue: IFAC-PapersOnLine, Lead Author: Thinh Doan, Year: 2019
 91. Title: Convolutional Graph Neural Networks, Venue: Asilomar SSC 2019, Lead Author: Fernando Gama, Year: 2019
 92. Title: Decentralization of Multiagent Policies by Learning What to Communicate, Venue: ICRA 2019, Lead Author: James Paulos, Year: 2019
 93. Title: DEDUCE: Diverse scEne Detection methods in Unseen Challenging Environments, Venue: IEEE/RSJ International Conference on Intelligent Robots and Systems, Lead Author: Anwesha Pal, Year: 2019
 94. Title: Differentiable Gaussian Process Motion Planning, Venue: ICRA, Lead Author: Mohak Bhardwaj, Year: 2020
 95. Title: Experimental Evaluation of Group Communications Protocols for Data Dissemination at the Tactical Edge, Venue: ICMCIS 2019, Lead Author: Niranjan Suri, Year: 2019

96. Title: Gated Graph Convolutional Recurrent Neural Networks, Venue: Eusipco 2019, Lead Author: Luana Ruiz, Year: 2019
97. Title: Generalizing Graph Convolutional Neural Networks with Edge-Variant Recursions on Graphs, Venue: Eusipco 2019, Lead Author: Elvin Isufi, Year: 2019
98. Title: Graph Embedding for the Division of Robotic Swarms, Venue: International Conference on Robotics and Automation (ICRA) Workshop on Robot Teammates Operating in Dynamic, Unstructured Environments (RT-DUNE), Lead Author: B. Reily, Year: 2019
99. Title: Human Sensitivity to Community Structure is Robust to Topological Variation, Venue: Complexity, Lead Author: Elizabeth Karuza, Year: 2019
100. Title: Hypothesis Assignment and Partial Likelihood Averaging for Cooperative Estimation, Venue: IEEE Conference on Decision and Control (CDC), Lead Author: Parth Paritosh, Year: 2019
101. Title: Information Filter Occupancy Mapping using Decomposable Radial Kernels, Venue: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Lead Author: Siwei Guo, Year: 2019
102. Title: Leveraging Experience in Lazy Search, Venue: Proceedings of Robotics Science and Systems XV (RSS), Lead Author: Mohak Bhardwaj, Year: 2019
103. Title: Linear Two-Time-Scale Stochastic Approximation: A Finite-Time Analysis, Venue: Allerton, Lead Author: Thinh Doan, Year: 2019
104. Title: Localization and Mapping using Instance-Specific Mesh Models, Venue: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Lead Author: Qiaojun Feng, Year: 2019
105. Title: Optimal Stochastic Vehicle Path Planning Using Covariance Steering, Venue: International Conference on Robotics and Automation, Lead Author: Kazuhide Okamoto, Year: 2019
106. Title: Optimal WDM Power Allocation via Deep Learning for Radio on Free Space Optics Systems, Venue: Globecom 2019, Lead Author: Zhan Gao, Year: 2019
107. Title: Perception-Aware Trajectory Generation for Aggressive Quadrotor Flight Using Differential Flatness, Venue: American Control Conference (ACC), Lead Author: Murali, Varun, Year: 2019

108. Title: Resilient Active Target Tracking with Multiple Robots, Venue: IEEE Robotics and Automation Letters, Lead Author: L. Zhou, Year: 2019
109. Title: Search and rescue under the forest canopy using multiple UAVs, Venue: IJRR, Lead Author: Yulun Tian, Year: 2020
110. Title: Sparse Learning of Parsimonious Reproducing Kernel Hilbert Space Models, Venue: ICASSP, Lead Author: Maria Peifer, Year: 2019
111. Title: The Blackbird UAV Dataset, Venue: IJRR, Lead Author: Amado Antonini, Year: 2020
112. Title: Unsupervised Role Discovery Using Temporal Observations of Agents, Venue: International Conference on Autonomous Agents and MultiAgent Systems, Lead Author: Andrew Silva, Year: 2019
113. Title: Visual Planning with Semi-Supervised Stochastic Action Representations, Venue: ICML 2019 Workshop Visual planning with semi-supervised stochastic action representations, Lead Author: Karl Schmeckpeper, Year: 2019
114. Title: WaveToFly: Using Gesture Commands to Direct UAVs, Venue: ICRA-WS RT_Dune, Lead Author: Shixin Li, Year: 2019
115. Title: Ad hoc Teamwork with Behavior Switching Agents, Venue: Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI), Lead Author: Manish Ravula, Year: 2019
116. Title: Approximated Dynamic Traits for Task Assignment in Heterogeneous Multi-Robot Teams, Venue: IEEE International Conference on Intelligent Robots and Systems (IROS), Lead Author: Glen Neville, Year: 2020
117. Title: Assumed Density Filtering Q Learning, Venue: IJCAI 2019, Lead Author: heejin Cloe Jeong, Year: 2019
118. Title: Building Self-Play Curricula Online by Playing with Expert Agents in Adversarial Games, Venue: Proceedings of the 8th Brazilian Conference on Intelligent Systems (BRACIS), Lead Author: Felipe Leno Da Silva, Year: 2019
119. Title: Desiderata for Planning Systems in General-Purpose Service Robots, Venue: Proceedings of the ICAPS Workshop on Planning and Robotics (PlanRob 2019), Lead Author: Nick Walker, Year: 2019

120. Title: Generative Adversarial Imitation from Observation, Venue: Imitation, Intent, and Interaction (I3) Workshop at ICML 2019, Lead Author: Faraz Torabi, Year: 2019
121. Title: Human Gaze-Driven Spatial Tasking of an Autonomous MAV, Venue: IEEE RA-L Robotics and Automation Letters and ICRA, 2019, Lead Author: Liangzhe Yuan, Year: 2019
122. Title: Imitation Learning from Video by Leveraging Proprioception, Venue: Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI), Lead Author: Faraz Torabi, Year: 2019
123. Title: Importance Sampling Policy Evaluation with an Estimated Behavior Policy, Venue: Proceedings of the 36th International Conference on Machine Learning (ICML), Lead Author: Josiah Hanna, Year: 2019
124. Title: Inertial Yaw-Independent Velocity and Attitude Estimation for High Speed Quadrotor Flight, Venue: IEEE RA-L Robotics and Automation Letters and ICRA, 2019, Lead Author: James Svacha, Year: 2019
125. Title: Input Hard Constrained Optimal Covariance Steering, Venue: CDC 2019, Lead Author: Kazuhide Okamoto, Year: 2019
126. Title: Lessons Learned From Deploying Autonomous Vehicles at UC San Diego, Venue: Field & Service Robotics, Lead Author: D. Paz, Year: 2019
127. Title: Leveraging Human Guidance for Deep Reinforcement Learning Tasks, Venue: Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI), Lead Author: Ruohan Zhang, Year: 2019
128. Title: Millimeter Wave Remote UAV Control and Communications for Public Safety Scenarios, Venue: International Workshop on Internet of Autonomous Unmanned Vehicles, IAUUV 2019, Lead Author: William Xia, Year: 2019
129. Title: Modeling mmWave Channels in High-Fidelity Simulations of Unmanned Aerial Systems, Venue: Signal Processing Advances in Wireless Communications (SPAWC), Lead Author: Tanmay Ram Godbole, Year: 2019
130. Title: Nonlinear Uncertainty Control with Iterative Covariance Steering, Venue: CDC 2019, Lead Author: Jack Ridderhof, Year: 2019
131. Title: Online Estimation of Geometric and Inertia Parameters for Multirotor Aerial Vehicles, Venue: IEEE International Conference on Robotics and Automation (ICRA) 2019, Lead Author: Valentin Wuest, Year: 2019

132. Title: Open-World Reasoning for Service Robots, Venue: Proceedings of the 29th International Conference on Automated Planning and Scheduling (ICAPS 2019), Lead Author: Yuqian Jiang, Year: 2019
133. Title: Optimal Temporal Logic Planning for Multi-Robot Systems in Uncertain Semantic Maps., Venue: IROS 2019, Lead Author: Ioannis Kantaros, Year: 2019
134. Title: Persistification of Robotic Tasks, Venue: IEEE Transactions on Control Systems Technology, Lead Author: Gennaro Notomista, Year: 2019
135. Title: Policy Improvement Directions for Reinforcement Learning in Reproducing Kernel Hilbert Spaces, Venue: Conference on Decision and Control, Lead Author: Santiago Paternain, Year: 2019
136. Title: Primal–Dual Gradient Dynamics for Cooperative Unknown Payload Manipulation Without Communication, Venue: ACC, Lead Author: Tatsuya Miyano, Year: 2020
137. Title: Recent Advances in Imitation Learning from Observation, Venue: Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI), Lead Author: Faraz Torabi, Year: 2019
138. Title: Resilience by Reconfiguration: Exploiting Heterogeneity in Robot Teams, Venue: IROS, Lead Author: Ragesh Ramachandran, Year: 2019
139. Title: Resilient Assignment Using Redundant Robots on Transport Networks with Uncertain Travel Time, Venue: IEEE T-ASE, Lead Author: Amanda Prorok, Year: 2019
140. Title: RIDM: Reinforced Inverse Dynamics Modeling for Learning from a Single Observed Demonstration, Venue: Imitation, Intent, and Interaction (I3) Workshop at ICML 2019, Lead Author: Brahma S. Pavse, Year: 2019
141. Title: Sample-Efficient Adversarial Imitation Learning from Observation, Venue: Imitation, Intent, and Interaction (I3) Workshop at ICML 2019, Lead Author: Faraz Torabi, Year: 2019
142. Title: Scalable Representation Learning for Long-Term Augmented Reality-Based Information Delivery In Collaborative Human-Robot Perception, Venue: International Conference on Virtual, Augmented and Mixed Reality (VAMR), Lead Author: Fei Han, Year: 2019
143. Title: Second-Order Filtering Algorithms for Streaming Optimization Problems, Venue: CAMSAP 2019, Lead Author: Tomer Harari Hamam, Year: 2019

144. Title: Solving Service Robot Tasks: UT Austin Villa@Home 2019 Team Report, Venue: AAAI Fall Symposium on Artificial Intelligence and Human-Robot Interaction for Service Robots in Human Environments (AI-HRI 2019), Lead Author: Rishi Shah, Year: 2019
145. Title: Stable, Concurrent Controller Composition for Multi-Objective Robotic Tasks, Venue: Proceedings of the 58th Conference on Decision and Control (CDC-2019), Lead Author: Anqi Li, Year: 2019
146. Title: Stochastic Latent Actor-Critic, Venue: NeurIPS, Lead Author: Sergey Levine, Year: 2020
147. Title: Task-Motion Planning with Reinforcement Learning for Adaptable Mobile Service Robots, Venue: Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2019), Lead Author: Yuqian Jiang, Year: 2019
148. Title: "Looking at the Right Stuff" - Guided Semantic-Gaze for Autonomous Driving, Venue: IEEE/CVF Conference on Computer Vision and Pattern Recognition, Lead Author: Anwesan Pal, Year: 2020
149. Title: Adaptive Task Allocation for Heterogeneous Multi-Robot Teams with Evolving and Unknown Robot Capabilities, Venue: IEEE International Conference on Robotics and Automation, Lead Author: Yousef Emam, Year: 2020
150. Title: CLEAR: A Consistent Lifting, Embedding, and Alignment Rectification Algorithm for Multi-View Data Association, Venue: Robotics: Science and Systems, Lead Author: Kaveh Fathian, Year: 2019
151. Title: Counterfactual Programming for Optimal Control, Venue: L4DC, Lead Author: Luiz Chamon, Year: 2020
152. Title: Controller Synthesis for Infinitesimally Shape-Similar Formations, Venue: IEEE International Conference on Robotics and Automation, Lead Author: Ian Buckley, Year: 2020
153. Title: DC-CAPT: Concurrent Assignment and Planning of Trajectories for Dubins Cars, Venue: IEEE International Conference on Robotics and Automation (ICRA), Lead Author: Michael Whitzer, Year: 2019
154. Title: Graph Policy Gradients for Large Scale Robot Control, Venue: CORL 2019, Lead Author: Arbaaz Khan, Year: 2019
155. Title: Human Information Processing in Complex Networks, Venue: Nature Physics, Lead Author: Christopher Lynn, Year: 2019

156. Title: Kimera: An Open-Source Library for Real-Time Metric-Semantic Localization and Mapping, Venue: ICRA, Lead Author: Antoni Rosinol, Year: 2020
157. Title: Learning Q-network for Active Information Acquisition, Venue: IEEE IROS, Lead Author: Heejin Jeong, Year: 2019
158. Title: Maximum Information Bounds for Planning Active Sensing Trajectories, Venue: IEEE IROS, Lead Author: Brent Schlotdfelt, Year: 2019
159. Title: Modular Robot Formation and Routing for Resilient Consensus, Venue: IEEE American Control Conference (ACC), Lead Author: Xi Yu, Year: 2019
160. Title: Multi-Agent Task Allocation using Cross-Entropy Temporal Logic Optimization, Venue: ICRA 2020, Lead Author: Christopher Banks, Year: 2020
161. Title: Multi-Robot Coordination for Estimation and Coverage of Unknown Spatial Fields, Venue: ICRA 2020, Lead Author: Alessia Benevento, Year: 2019
162. Title: Multi-Robot Path Deconfliction through Prioritization by Path Prospects, Venue: IEEE R-AL, Lead Author: Wenying Wu, Year: 2020
163. Title: Optimal Computation-Communication Trade-offs in Processing Networks, Venue: IEEE Transactions on Network Science and Engineering, Lead Author: Luca Ballotta, Year: 2020
164. Title: Optimization-Based Distributed Flocking Control for Multiple Rigid Bodies, Venue: IEEE Robotics and Automation Letters, Lead Author: Tatsuya Ibuki, Year: 2020
165. Title: Planning with Uncertain Specifications (PUnS), Venue: IEEE Robotics and Automation Letters, Lead Author: Ankit Shah, Year: 2020
166. Title: Representing Multi-Robot Structure through Multimodal Graph Embedding for the Selection of Robot Teams, Venue: 2020 IEEE International Conference on Robotics and Automation (ICRA), Lead Author: Brian Reily, Year: 2020
167. Title: Semi-Supervised Learning of Decision-Making Models for Human-Robot Collaboration, Venue: CORL, Lead Author: Vaibhav Unhelkar, Year: 2019

168. Title: Simultaneous Learning from Human Pose and Object Cues for Real-Time Activity Recognition, Venue: 2020 IEEE International Conference on Robotics and Automation (ICRA), Lead Author: Brian Reily, Year: 2020
169. Title: Target Driven Visual Navigation Exploiting Object Relationships, Venue: IEEE Transactions on Robotics, Lead Author: Kaveh Fathian, Year: 2020
170. Title: Team Composition for Perimeter Defense with Patrollers and Defenders, Venue: 58th IEEE Conference on Decision and Control (CDC 2019), Lead Author: Daigo Shishika, Year: 2019
171. Title: TEASER: Fast and Certifiable Point Cloud Registration, Venue: IEEE Transactions on Robotics (TRO), Lead Author: Heng Yang, Year: 2020
172. Title: Towards Online Observability-Aware Trajectory Optimization for Landmark-Based Estimators, Venue: <https://arxiv.org/abs/1908.03790>, Lead Author: Kris Frey, Year: 2019
173. Title: Visual Coverage Maintenance for Quadcopters Using Nonsmooth Barrier Functions, Venue: ICRA 2020, Lead Author: Riku Funada, Year: 2020
174. Title: “Looking at the right stuff” - Guided semantic-gaze for autonomous driving,” Venue: Computer Vision and Pattern Recognition (CVPR), Lead Author: Anwesan Pal, Year: 2020
175. Title: 3D Dynamic Scene Graphs: Actionable Spatial Perception with Places, Objects, and Humans, Venue: RSS, Lead Author: Antoni Rosinol, Year: 2020
176. Title: A Polynomial-Time Solution for Robust Registration with Extreme Outlier Rates, Venue: Robotics: Science and Systems (RSS), Lead Author: Heng Yang, Year: 2019
177. Title: A Quaternion-Based Certifiably Optimal Solution to the Wahba Problem with Outliers, Venue: International Conference on Computer Vision (ICCV), Lead Author: Heng Yang, Year: 2019
178. Title: A Zeroth-order Learning Algorithm for Ergodic Optimization of Wireless Systems with No Models and No Gradients, Venue: 45th IEEE International Conference on Acoustics, Speech and Signal Processing, Lead Author: Dionysios Kalogerias, Year: 2020
179. Title: Almost-Zero Duality Gaps in Model-Free Resource Allocation for Wireless Systems, Venue: EUSIPCO, Lead Author: Dionysios Kalogerias, Year: 2020

180. Title: Approximate Supermodularity of Kalman Filter Sensor Selection, Venue: IEEE Transactions on Automatic Control, Lead Author: Luiz Chamon, Year: 2020
181. Title: Architecture and Evolution of Semantic Networks in Mathematics Texts, Venue: Proceedings of the Royal Society A, Lead Author: Nicolas Christianson, Year: 2020
182. Title: Autonomous Navigation in Unknown Environments using Sparse Kernel-based Occupancy Mapping, Venue: IEEE International Conference on Robotics and Automation (ICRA), Lead Author: Thai Duong, Year: 2020
183. Title: Cooperative Team Strategies for Multi-Player Perimeter-Defense Games, Venue: IEEE Robotics and Automation Letters, Lead Author: Daigo Shishika, Year: 2020
184. Title: Counterfactual Programming for Optimal Control, Venue: 2nd Conference on Learning for Dynamics and Control, Lead Author: Luiz Chamon, Year: 2020
185. Title: Covariance Steering for Discrete-Time Linear-Quadratic Stochastic Dynamic Games, Venue: IEEE Conference on Decision and Control (CDC), Lead Author: Ramana Makkapati, Year: 2020
186. Title: Decentralized Minimum-Energy Coverage Control for Time-Varying Density Functions, Venue: International Symposium on Multi-Robot and Multi-Agent Systems (MRS), Lead Author: Maria Santos, Year: 2019
187. Title: Deep Imitative Models for Flexible Inference, Planning, and Control, Venue: International Conference on Learning Representations (ICLR), Lead Author: Nicholas Rhinehart, Year: 2020
188. Title: Dense r-Robust Formations on Lattices, Venue: IEEE International Conference on Robotics and Automation, Lead Author: Luis Guerrero, Year: 2020
189. Title: DOOR-SLAM: Distributed, Online, and Outlier Resilient Slam for Robotic Teams, Venue: IEEE Robotics and Automation Letters (RA-L), Lead Author: Pierre-Yves Lajoie, Year: 2020
190. Title: Dynamic Target Tracking and Energy Efficient AUV Path Planning for Trash Collection Using Ocean Current, Venue: Oceans, Lead Author: Michelle Sit, Year: 2020
191. Title: EdgeNets: Edge Varying Graph Neural Networks, Venue: TPAMI, Lead Author: Elvin Isufi, Year: 2021

192. Title: Fast and Safe Path-Following Control using a State-Dependent Directional Metric, Venue: IEEE International Conference on Robotics and Automation (ICRA), Lead Author: Zhichao Li, Year: 2020
193. Title: Federated Classification with low Complexity Reproducing Kernel Hilbert Space Representations, Venue: 45th IEEE International Conference on Acoustics, Speech and Signal Processing, Lead Author: Peifer Maria, Year: 2020
194. Title: From Sensor to Processing Networks: Optimal Estimation with Computation and Communication Latency, Venue: IFAC World Congress, Lead Author: Luca Ballotta, Year: 2020
195. Title: Functional Brain Network Architecture Supporting the Learning of Social Networks in Humans, Venue: NeuroImage, Lead Author: Steve Thompson, Year: 2020
196. Title: Functional Nonlinear Sparse Models, Venue: IEEE Transactions on Signal Processing, Lead Author: Luiz Chamon, Year: 2020
197. Title: Game Theoretic Formation Design for Probabilistic Barrier Coverage, Venue: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Lead Author: Daigo Shishika, Year: 2020
198. Title: Gated Graph Recurrent Neural Networks, Venue: IEEE Transactions on Signal Processing, Lead Author: Luana Ruiz, Year: 2020
199. Title: Graduated Non-convexity for Robust Spatial Perception: From non-minimal solvers to global outlier rejection, Venue: IEEE Robotics and Automation Letters (RA-L), Lead Author: Heng Yang, Year: 2020
200. Title: Graph Learning: Inferring the Network Structure of the Environment, Venue: Proceedings of the National Academy of the Sciences, Lead Author: Christophe W. Lynn, Year: 2020
201. Title: Graph Neural Networks for Decentralized Multi-Robot Path Planning, Venue: IROS, Lead Author: Qingbiao Li, Year:
202. Title: Graph, Convolutions and Neural Networks, Venue: IEEE Signal Processing Magazine, Lead Author: Fernando Gama, Year: 2020
203. Title: Graphon Filters: Signal Processing in Very Large Graphs, Venue: EUSIPCO, Lead Author: Luana Ruiz, Year: 2021
204. Title: Graphon Pooling in Graph Neural Networks, Venue: EUSIPCO, Lead Author: Alejandro Parada Mayoraga, Year: 2021

205. Title: Graphon Signal Processing, Venue: IEEE Transactions on Signal Processing, Lead Author: Luana Ruiz, Year: 2021
206. Title: In Perfect Shape: Certifiably Optimal 3D Shape Reconstruction from 2D Landmarks, Venue: IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), Lead Author: Heng Yang, Year: 2020
207. Title: Individual Differences in Learning Social and Nonsocial Network Structures, Venue: J Exp Psychol Learn Mem Cogn., Lead Author: Steve Thompson, Year: 2019
208. Title: Invariance-Preserving Localized Activation Functions for Graph Neural Networks, Venue: IEEE Transactions on Signal Processing, Lead Author: Luana Ruiz, Year: 2020
209. Title: Invertible Generalized Synchronization: A Putative Mechanism for Implicit Learning in Biological and Artificial Neural Systems, Venue: Chaos, Lead Author: Zhixin Lu, Year: 2020
210. Title: Latent State Models for Meta-Reinforcement Learning, Venue: CoRL, Lead Author: Anusha Nagabandi, Year: 2020
211. Title: Learning Constrained Resource Allocation Policies in Wireless Control Systems, Venue: In 59th IEEE Conference on Decision and Control, Lead Author: Vinicius Lima, Year: 2020
212. Title: Learning Hierarchical Relationships for Object-Goal Navigation, Venue: Conference on Robot Learning, Lead Author: Cassie Y. Qiu, Year: 2020
213. Title: Learning Navigation Costs from Demonstration in Partially Observable Environments, Venue: IEEE International Conference on Robotics and Automation (ICRA), Lead Author: Tianyu Wang, Year: 2020
214. Title: Learning Navigation Costs from Demonstrations with Semantic Observations, Venue: Learning for Dynamics and Control (L4DC), Lead Author: Tianyu Wang, Year: 2020
215. Title: Learning Predictive Models From Observation and Interaction, Venue: European Conference on Computer Vision, Lead Author: Karl Schmeckpepper, Year: 2020
216. Title: Mobile Wireless Network Infrastructure on Demand, Venue: International Conference on Robotics and Automation (ICRA), Lead Author: Daniel Mox, Year: 2020

217. Title: Model-Based Meta-Reinforcement Learning for Flight with Suspended Payloads, Venue: ICRA, Lead Author: Suneel Belkhale, Year: 2021
218. Title: Multi-mode Autonomous Communication Systems, Venue: Asilomar Conference on Signals, Systems, and Computers,, Lead Author: Miguel Calvo-Fullana, Year: 2019
219. Title: Navigation of a Quadratic Potential with Star Obstacles, Venue: 2020 IEEE American Control Conference, Lead Author: Harshat Kumar, Year: 2020
220. Title: Network Architectures Supporting Learnability, Venue: Proceedings of the Royal Society B, Lead Author: Perry Zurn, Year: 2020
221. Title: Optimal Algorithms for Submodular Maximization with Distributed Constraints., Venue: Learning for Dynamics and Control, Lead Author: Alex Robey, Year: 2021
222. Title: Optimal Power Flow Using Graph Neural Networks, Venue: 45th IEEE International Conference on Acoustics, Speech and Signal Processing, Lead Author: Damian Owerko, Year: 2020
223. Title: Outlier-Robust Spatial Perception: Hardness, General-Purpose Algorithms, and Guarantees, Venue: IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS), Lead Author: Vasileios Tzoumas, Year: 2019
224. Title: Perimeter-Defense Game Between Aerial Defender and Ground Intruder, Venue: 59th IEEE Conference on Decision and Control (CDC), Lead Author: Elijah S. Lee, Year: 2020
225. Title: Probabilistic Safety Constraints for Learned High Relative-Degree System Dynamics, Venue: Learning for Dynamics and Control (L4DC), Lead Author: Mohammad Javad Khojasteh, Year: 2020
226. Title: Realization of r-Robust Formations in the Plane Using Control Barrier Functions, Venue: IEEE Control Systems Letters, Lead Author: Luis Guerrero, Year: 2020
227. Title: Resilient Control: Compromising to Adapt, Venue: IEEE Conference on Decision and Control, Lead Author: Luiz Chamon, Year: 2020
228. Title: Resilient Coverage: Exploring the Local-to-Global Trade-off, Venue: IROS, Lead Author: Ragesh Ramachandran, Year: 2021
229. Title: Resilient Information Acquisition for Multi Robot Teams, Venue: IEEE Transactions on Robotics, Lead Author: Brent Schlotfeldt, Year: 2021

230. Title: Resource Allocation in Large-Scale Wireless Control Systems with Graph Neural Networks, Venue: 21st IFAC World Congress, Lead Author: Vinicius Lima Silva, Year: 2021
231. Title: Resource Allocation in Wireless Control Systems via Deep Policy Gradient, Venue: SPAWC 2020, Lead Author: Vinicius Lima, Year: 2020
232. Title: Risk-Constrained Linear-Quadratic Regulators, Venue: IEEE Conference on Decision and Control, Lead Author: Anastasios Tsiamis, Year: 2020
233. Title: Robust and Efficient Forward, Differential, and Inverse Kinematics using Dual Quaternions, Venue: IJRR, Lead Author: Neil Dantam, Year: 2020
234. Title: Robust Assignment Using Redundant Robots on Transport Networks with Uncertain Travel Time, Venue: IEEE Transactions on Automation Science and Engineering (T-ASE), Lead Author: Amanda Prorok, Year: 2020
235. Title: Scaling Up Multiagent Reinforcement Learning for Robotic Systems: Learn an Adaptive Sparse Communication Graph, Venue: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Lead Author: Chuangchuang Sun, Year: 2020
236. Title: Sparse Multiresolution Representation with Adaptive Kernels, Venue: IEEE Transactions on Signal Processing, Lead Author: Maria Peifer, Year: 2020
237. Title: Spatial Gating Strategies for Graph Recurrent Neural Networks, Venue: 45th IEEE International Conference on Acoustics, Speech and Signal Processing, Lead Author: Luana Ruiz, Year: 2020
238. Title: Stability of Graph Neural Networks to Relative Perturbations, Venue: 45th IEEE International Conference on Acoustics, Speech and Signal Processing, Lead Author: Fernando Gama, Year: 2020
239. Title: Stochastic Policy Gradient Ascent in Reproducing Kernel Hilbert Spaces, Venue: IEEE Transactions on Automatic Control, Lead Author: Santiago Paternain, Year: 2020
240. Title: Structure from Noise: Mental Errors Yield Abstract Representations of Events, Venue: Nature Communications, Lead Author: Christopher Lynn, Year: 2020
241. Title: Sufficiently Accurate Model Learning, Venue: 2020 International Conference on Robotics and Automation, Lead Author: Clark Zhang, Year: 2020

242. Title: Supervised Chaotic Source Separation by a Tank of Water, Venue: Chaos, Lead Author: Zhixin Lu, Year: 2020
243. Title: Synthesis of a Time-Varying Communication Network by Robot Teams With Information Propagation Guarantees, Venue: IEEE Robotics and Automation Letters, Lead Author: Xi Yu, Year: 2020
244. Title: The Empirical Duality Gap of Constrained Statistical Learning, Venue: 45th IEEE International Conference on Acoustics, Speech and Signal Processing, Lead Author: Luiz Chamon, Year: 2020
245. Title: The Graphon Fourier Transform, Venue: 45th IEEE International Conference on Acoustics, Speech and Signal Processing, Lead Author: Luana Ruiz, Year: 2020
246. Title: A Distributed Pipeline for Scalable, Deconflicted Formation Flying, Venue: IEEE Robotics and Automation Letters (RA-L), Lead Author: Parker Lusk, Year: 2020
247. Title: Asynchronous and Parallel Distributed Pose Graph Optimization, Venue: IEEE Robotics and Automation Letters (RA-L), Lead Author: Yulun Tian, Year: 2020
248. Title: Better Safe Than Sorry: Risk-Aware Nonlinear Bayesian Estimation, Venue: 45th IEEE International Conference on Acoustics, Speech and Signal Processing, Lead Author: Dionysios Kalogerias, Year: 2020
249. Title: Chance-Constrained Optimal Covariance Steering with Iterative Risk Allocation, Venue: American Control Conference, Lead Author: Joshua Pilipovksy, Year: 2020
250. Title: Computation-Communication Trade-offs and Sensor Selection in Real-time Estimation for Processing Networks, Venue: IEEE Trans. on Network Science and Engineering, Lead Author: Luca Ballotta, Year: 2020
251. Title: Consensus of Multi-Agent Systems with Asynchronous Cloud-Communication, Venue: Automatica, Lead Author: Sean Bowman, Year: 2020
252. Title: Covariance Steering with Optimal Risk Allocation, Venue: IEEE Transactions on Aerospace and Electronic Systems, Lead Author: Joshua Pilipovksy, Year: 2020
253. Title: Dense Incremental Metric-Semantic Mapping via Sparse Gaussian Process Regression, Venue: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Lead Author: Ehsan Zobeidi, Year: 2020

254. Title: Distributed Attack-Robust Submodular Maximization for Multi-Robot Planning, Venue: ICRA, Lead Author: Lifeng Zhou, Year: 2020
255. Title: Finite Time Performance of Distributed Two-Time Scale Stochastic Approximation, Venue: ACC 2020, Lead Author: Think Doan, Year: 2020
256. Title: Forward Chaining Hierarchical Partial-Order Planning, Venue: Workshop on the Algorithmic Foundations of Robotics, Lead Author: Andrew Messing, Year: 2020
257. Title: Fully Convolutional Geometric Features for Category-level Object Alignment, Venue: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Lead Author: Qiaojun Feng, Year: 2020
258. Title: Infinitesimal Shape-Similarity for Characterization and Control of Bearing-Only Multi-Robot Formations, Venue: IEEE Transactions on Robotics, Lead Author: Ian Buckley, Year: 2020
259. Title: Information Theoretic Active Exploration in Signed Distance Fields, Venue: ICRA, Lead Author: Kelsey Saulnier, Year: 2020
260. Title: Learned Sampling Distributions for Efficient Planning in Hybrid Geometric and Object-Level Representations., Venue: International Conference on Robotics and Automation, Lead Author: Martina Stadler, Year: 2020
261. Title: Learning Predictive Models From Observation and Interaction, Venue: European Conference on Computer Vision, Lead Author: Karl Schmeckpeper, Year: 2020
262. Title: Marginal Density Averaging for Distributed Node Localization from Local Edge Measurements, Venue: IEEE Conference on Decision and Control (CDC), Lead Author: Parth Paritosh, Year: 2020
263. Title: Metrically-Scaled Monocular SLAM using Learned Scale Factors, Venue: 2020 IEEE International Conference on Robotics and Automation (ICRA), Lead Author: Nick Greene, Year: 2020
264. Title: One Ring to Rule Them All: Certifiably Robust Geometric Perception with Outliers, Venue: 34th Conference on Neural Information Processing Systems (NeurIPS 2020), Lead Author: Heng Yang, Year: 2020"
265. Title: OrcVIO: Object Residual Constrained Visual-Inertial Odometry, Venue: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Lead Author: Mo Shan, Year: 2020

266. Title: Reactive Temporal Logic Planning for Multiple Robots in Unknown Environments., Venue: ICRA, Lead Author: Yannis Kantaros, Year: 2020
267. Title: Resilience in Multi-Robot Target Tracking Through Reconfiguration, Venue: ICRA, Lead Author: Ragesh Ramachandran, Year: 2020
268. Title: Semantic Trajectory Planning for Long-Distant Unmanned Aerial Vehicle Navigation in Urban Environments., Venue: International Conference on Intelligent Robots and Systems (IROS), Lead Author: Markus Ryll, Year: 2021
269. Title: Steering the State of Linear Stochastic Systems: A Constrained Minimum Principle Formulation, Venue: American Control Conference, Lead Author: Ali Pakniyat, Year: 2020
270. Title: STRATA: A Unified Framework for Task Assignments in Large Teams of Heterogeneous Robots, Venue: Journal of Autonomous Agents and Multi-Agent Systems (J-AAMAS), Lead Author: Harish Ravichandar, Year: 2020
271. Title: Tracking and Relative Localization of Drone Swarms with a Vision-based Headset, Venue: IEEE Robotics and Automation Letters, Lead Author: Maxim Pavliv, Year: 2021
272. Title: A Continuous Representation of Belief over $SO(3)$ for Robust Rotation Learning with Uncertainty, Venue: Robotics: Science and Systems (RSS), Lead Author: Valentin Peretroukhin, Year: 2020
273. Title: A Policy Gradient Algorithm for Learning to Learn in Multiagent Reinforcement Learning, Venue: ICML 2021, Lead Author: Dong-Ki Kim, Year: 2021
274. Title: A Resource-Aware Approach to Collaborative Loop Closure Detection with Provable Performance Guarantees, Venue: International Journal of Robotics Research (IJRR), Lead Author: Yulun Tian, Year: 2020
275. Title: Adaptive Partitioning for Cooperative Multi-agent Perimeter Defense., Venue: IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), Lead Author: Douglas Guimares Macharet, Year: 2020
276. Title: Approximated Dynamic Trait Models for Heterogeneous Multi-Robot Teams, Venue: IROS 2020, Lead Author: Glen Neville, Year: 2020
277. Title: Asynchronous Adaptive Sampling and Reduced-Order Modeling of Dynamic Processes by Robot Teams via Intermittently Connected Networks, Venue: International Conference on Intelligent Robots and Systems (IROS), Lead Author: Hannes Rovina, Year: 2020

278. Title: Autonomous Vehicle Benchmarking Using Unbiased Metrics, Venue: International Conference on Intelligent Robots and Systems (IROS), Lead Author: David Paz-Ruiz, Year: 2020
279. Title: Distributed Attack-Submodular Submodular Maximization, Venue: IEEE Transaction on Robotics, Lead Author: Lifeng Zhou, Year: 2020
280. Title: FENet: Fast Real-time Semantic Edge Detection Network, Venue: IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR), Lead Author: Yang Zhou, Year: 2020
281. Title: Finite Sample Analysis of a Two-Time-Scale Greedy Natural Actor-Critic Algorithm, Venue: IEEE Trans. Automatic Control, Lead Author: Sajad Khodadadian, Year: 2021
282. Title: FISAR: Forward Invariant Safe Reinforcement Learning with a Deep Neural Network-Based Optimize, Venue: ICML 2020, ICRA 2021, Lead Author: Chuangchuang Sun, Year: 2020
283. Title: Interactive Multi-Robot Painting Through Colored Motion Trails, Venue: Frontiers in Robotics and AI, Lead Author: Maria Santos, Year: 2020
284. Title: Long-Horizon Planning with Goal-Conditioned Hierarchical Prediction, Venue: NeurIPS, Lead Author: Oleh Rybkin, Year: 2020
285. Title: MAMS-A*: Multi-Agent Multi-Scale A*, Venue: International Conference on Robotics and Automation, Lead Author: Jaein Lim, Year: 2020
286. Title: Probabilistic Semantic Mapping for Urban Autonomous Driving Applications, Venue: International Conference on Intelligent Robots and Systems (IROS), Lead Author: David Paz Ruiz, Year: 2020
287. Title: Q-Tree Search: An Information-Theoretic Approach Towards Hierarchical Abstractions for Agents with Computational Limitations, Venue: IEEE Transactions on Robotics, Lead Author: Daniel Larsson, Year: 2020
288. Title: Resilient Active Information Acquisition with Teams of Robots, Venue: IEEE Transaction on Robotics, Lead Author: Brent Schlotfeldt, Year: 2021
289. Title: Robust and Adaptive Sequential Submodular Optimization., Venue: IEEE Transaction on Automatic Control, Lead Author: Vasileios Tzoumas, Year: 2020"
290. Title: Target Driven Visual Navigation Exploiting Object Relationships, Venue: Conference on Robot Learning, Lead Author: Y. Cassie Qiu, Year: 2020

291. Title: Primal-Dual Mesh Convolutional Neural Networks, Venue: Conference on Neural Information Processing Systems, Venue: NeurIPS 2020, Lead Author: Francesco Milano, Year: 2020
292. Title: Resilience in Multi-Robot Multi-Target Tracking with unknown number of targets through reconfiguration, Venue: IEEE Transactions on Control of Network Systems, Lead Author: Ragesh Ramachandran, Year: 2021
293. Title: A Generalized A* Algorithm for Finding Globally Optimal Paths in Weighted Colored Graphs, Venue: ICRA 2021, Lead Author: Jaein Lim, Year: 2021
294. Title: A Resilient and Energy-Aware Task Allocation Framework for Heterogeneous Multi-Robot Systems, Venue: IEEE Transactions on Robotics, Lead Author: Gennaro Notomista, Year: 2021
295. Title: Active Bayesian Multi-class Mapping from Range and Semantic Segmentation Observations, Venue: IEEE ICRA 2021, Lead Author: Arash Asgharivaskasi, Year: 2021
296. Title: Any Way You Look At It: Semantic Crossview Localization and Mapping with LiDAR, Venue: IEEE RA-L , Lead Author: Ian D. Miller, Year: 2021
297. Title: Area Defense and Surveillance on Rectangular Regions Using Control Barrier Functions, Venue: IEEE IROS, Lead Author: Luis Guerrero Bonilla, Year: 2021
298. Title: Bounded-Rational Pursuit-Evasion Games, Venue: American Control Conference 2021, Lead Author: Scott Guan, Year: 2021
299. Title: Coding for Distributed Multi-Agent Reinforcement Learning, Venue: IEEE ICRA 2021, Lead Author: Baoqian Wang, Year: 2021
300. Title: Data-Driven Adaptive Task Allocation for Heterogeneous Multi-Robot Teams Using Robust Control Barrier Functions, Venue: IEEE International Conference on Robotics and Automation, Lead Author: Yousef Emam, Year: 2021
301. Title: Decentralized Wireless Resource Allocation with Graph Neural Networks, Venue: ASILOMAR 2020, Lead Author: Zhiyang Wang, Year: 2020
302. Title: Fair Robust Assignment using Redundancy, Venue: RA-L 2021, Lead Author: Matthew Malencia, Year: 2021

303. Title: Hierarchical Object Map Estimation for Efficient and Robust Navigation, Venue: ICRA 2021, Lead Author: Kyel Ok, Year: 2021
304. Title: How humans learn and represent networks, Venue: Proceedings of the National Academy of the Sciences, Lead Author: Christopher W. Lynn, Year: 2020
305. Title: Information-Theoretic Abstractions for Planning for Agents with Computational Constraints, Venue: IEEE RA-L, Lead Author: Daniel Larsson, Year: 2021
306. Title: Learning Connectivity-Maximizing Network Configurations, Venue: IEEE RA-L/ICRA, Lead Author: Daniel Mox, Year: 2022
307. Title: Learning Nash Equilibria in Zero-Sum Stochastic Games via Entropy-Regularized Policy Approximation, Venue: 30th International Joint Conference on Artificial Intelligence, Lead Author: Quifan Zhang, Year: 2021
308. Title: Message-Aware Graph Attention Networks for Large-Scale Multi-Robot Path Planning, Venue: IEEE R-AL, Lead Author: Qingbiao Li, Year: 2021
309. Title: MIXER: A Principled Framework for Multimodal, Multiway Data Association, Venue: workshop on Robust Perception For Autonomous Field Robots in Challenging Environments, IEEE ICRA, Lead Author: Parker Lusk, Year: 2021
310. Title: MultiViewStereoNet: Fast Multi-View Stereo Depth Estimation Using Incremental Viewpoint-Compensated Feature Extraction, Venue: ICRA 2021, Lead Author: Nick Greene, Year: 2021
311. Title: Non-Monotone Energy-Aware Information Gathering for Heterogeneous Robot Teams, Venue: IEEE ICRA 2021, Lead Author: Xiaoyi Cai, Year: 2021
312. Title: Optimal Controller Synthesis and Dynamic Quantizer Switching for Linear-Quadratic-Gaussian Systems, Venue: IEEE Transactions on Automatic Control, Lead Author: Dipankar Maity, Year: 2021
313. Title: Optimizing Non-Markovian Information Gain under Physics-based Communication Constraints, Venue: IEEE RA-L, Lead Author: Matthew Schack, Year: 2021
314. Title: Resilient Task Allocation in Heterogeneous Multi-Robot Systems, Venue: IEEE RA-L, Lead Author: Siddharth Mayya, Year: 2021

315. Title: Sampling-based Planning for Non-myopic Multi-Robot Information Gathering, Venue: Autonomous Robots, Lead Author: Yiannis Kantaros, Year: 2021
316. Title: Toward Robust and Efficient Online Adaptation for Deep Stereo Depth Estimation, Venue: ICRA 2021, Lead Author: Milo Knowles, Year: 2021
317. Title: Tractable Mechanisms for computing near-optimal utility functions, Venue: 20th International Conference on Autonomous Agents and Multiagent Systems, Lead Author: Rohit Konda, Year: 2021
318. Title: Bounded Rationality in Learning, Perception, Decision-Making, and Stochastic Games, Venue: Handbook of Reinforcement Learning and Control, Lead Author: Panagiotis Tsiotras, Year: 2021
319. Title: CLIPPER: A Graph-Theoretic Framework for Robust Data Association, Venue: IEEE International Conference on Robotics and Automation, Lead Author: Parker Lusk, Year: 2021
320. Title: Compressibility of complex networks, Venue: Proceedings of the National Academy of the Sciences, Lead Author: Christopher W. Lynn, Year: 2021
321. Title: Distributed Certifiably Correct Pose-Graph Optimization, Venue: IEEE Transactions on Robotics, Lead Author: Yulun Tian, Year: 2021
322. Title: Fast Convergence Rates of Distributed Subgradient Methods with Adaptive Quantization, Venue: IEEE TAC, Lead Author: Thinh Doan, Year: 2021
323. Title: Unsupervised Learning for Asynchronous Resource Allocation In Ad-Hoc Wireless Networks, Venue: ICASSP 2021, Lead Author: Zhiyang Wang, Year: 2021
324. Title: Defending a Perimeter from a Ground Intruder Using an Aerial Defender: Theory and Practice, Venue: IEEE SSRR 2022, Lead Author: Elijah S. Lee, Year: 2022
325. Title: A Decentralized Policy Gradient Approach to Multi-Task Reinforcement Learning, Venue: UAI, Lead Author: Sihan Zeng, Year: 2021
326. Title: Active Exploration and Mapping via Iterative Covariance Regulation over Continuous SE(3) Trajectories, Venue: IEEE/RSJ IROS 2021, Lead Author: Shumon Koga, Year: 2021

327. Title: An Interleaved Approach to Trait-Based Task Allocation and Scheduling, Venue: IROS 2021, Lead Author: Glen Neville, Year: 2021
328. Title: CORSAIR: Convolutional Object Retrieval and Symmetry-Aided Registration, Venue: IEEE/RSJ IROS 2021, Lead Author: Tianyu Zhao, Year: 2021
329. Title: Data-Driven Robust Barrier Functions for Safe, Long-Term Operation, Venue: IEEE Transactions on Robotics, Lead Author: Yousef Emam, Year: 2021
330. Title: Distributed Heterogeneous Multi-robot Source Seeking Using Information Based Sampling with Visual Recognition, Venue: ISER 2021, Lead Author: Carlos Nieto, Year: 2021
331. Title: Dynamic Defender-Attacker Blotto Game, Venue: IEEE ACC 2022, Lead Author: Daigo Shishika, Year: 2022
332. Title: ELLIPSDF: Joint Object Pose and Shape Optimization with a Bi-level Ellipsoid and Signed Distance Function Description, Venue: IEEE/CVF ICCV 2021, Lead Author: Mo Shan, Year: 2021
333. Title: Guarding a Translating Line with an Attached Defender, Venue: IEEE ACC 2022, Lead Author: Goutam Das, Year: 2022
334. Title: Kimera: from SLAM to Spatial Perception with 3D Dynamic Scene Graphs, Venue: IJRR, Lead Author: Antoni Rosinol, Year: 2021
335. Title: Multi-Agent Consensus Subject to Communication and Privacy Constraints, Venue: IEEE Transactions on Control of Network Systems, Lead Author: Dipankar Maity, Year: 2021
336. Title: Multi-Robot Coverage and Exploration using Spatial Graph Neural Networks, Venue: IROS 2021, Lead Author: Ekaterina Tolstaya, Year: 2021
337. Title: Optimal Pose and Shape Estimation for Category-Level 3D Object Perception, Venue: RSS, Lead Author: Jingnan Shi, Year: 2021
338. Title: Resilience in Multirobot Multitarget Tracking With Unknown Number of Targets Through Reconfiguration, Venue: IEEE Transactions on Control of Network Systems, Lead Author: Ragesh Ramachandran, Year: 2021
339. Title: Resilient Monitoring in Heterogeneous Multi-robot Systems through Network Reconfiguration, Venue: IEEE Transactions on Robotics, Lead Author: Rajesh Ramachandran, Year: 2021

340. Title: Shaping Large Population Agent Behaviors Through Entropy-Regularized Mean-Field Games, Venue: American Control Conference, Lead Author: Yue Guan, Year: 2021
341. Title: A Linear Programming Approach for Resource-Aware Information-Theoretic Tree Abstractions, Venue: Computation-aware Algorithmic Design for Cyber-Physical Systems, Lead Author: Daniel Larsson, Year: 2021
342. Title: Adaptive and Risk Aware Target Tracking with Heterogeneous Robot Teams, Venue: IEEE ICRA 2022 + RAL, Lead Author: Siddharth Mayya, Year: 2022
343. Title: Adaptive Wireless Power Allocation with Graph Neural Networks, Venue: IEEE ICASSP, Lead Author: Navid NaderiAlizadeh, Year: 2022
344. Title: An Agnostic Approach to Federated Learning with Class Imbalance, Venue: ICLR 2022, Lead Author: Zebang Shen, Year: 2022
345. Title: Convergence Rates of Distributed Gradient Methods Under Random Quantization: A Stochastic Approximation Approach, Venue: IEEE TAC, Lead Author: Thinh Doan, Year: 2021
346. Title: Coverage Control in Multi-Robot Systems via Graph Neural Networks, Venue: IEEE ICRA 2022, Lead Author: Walker Gosrich, Year: 2022
347. Title: Explaining Reward Functions to Humans for Better Human-Robot Collaboration, Venue: AI-HRI AAAI Fall Symposium, Lead Author: Lindsay Sanneman, Year: 2021
348. Title: Finite Sample Analysis of Two-Time-Scale Natural Actor-Critic Algorithm, Venue: IEEE TAC, Lead Author: Sajad Khodadadian, Year: 2021
349. Title: Finite-Time Analysis of Decentralized Stochastic Approximation with Applications in Multi-Agent and Multi-Task Learning, Venue: CDC 2022, Lead Author: S. Zeng, Year: 2022
350. Title: GRSTAPS: Graphically Recursive Simultaneous Task Allocation, Planning, and Scheduling, Venue: IJRR, Lead Author: Andrew Messing, Year: 2021
351. Title: Heterogeneous Multi-Robot Adversarial Patrolling Using Poly-Matrix Games, Venue: ICRAI 2021, Lead Author: Alex Langley, Year: 2021
352. Title: Infrastructure Node-based Vehicle Localization for Autonomous Driving, Venue: IEEE Intelligent Vehicle 2022, Lead Author: Elijah S. Lee, Year: 2022

353. Title: Kimera-Multi: Robust, Distributed, DenseMetric-Semantic SLAM for Multi-Robot Systems, Venue: TRO, Lead Author: Yulun Tian, Year: 2021"
354. Title: Monitoring and diagnosability of perception systems., Venue: IROS 2021, Lead Author: Pasquale Antonante, Year: 2021
355. Title: Multi-Task Reinforcement Learning in Reproducing Kernel Hilbert Spaces via Cross-Learning, Venue: TSP , Lead Author: Juan Cervino, Year: 2021
356. Title: Neural Trees for Learning on Graphs. In Conference on Neural Information Processing Systems, Venue: Neurips, Lead Author: Rajat Talak, Year: 2021
357. Title: Optimization-Based Robot Team Exploration Considering Attrition and Communication Constraints, Venue: IROS 2021, Lead Author: Matthew Schack, Year: 2021
358. Title: Resilient Coalition Formation in Heterogeneous Teams via Imitation Learning, Venue: IEEE IROS2021 workshop, Lead Author: Anusha Srikanthan, Year: 2022
359. Title: Resource-Aware Adaptation of Heterogeneous Strategies for Coalition Formation, Venue: AAMAS 2022 Extended Abstract, Lead Author: Anusha Srikanthan, Year: 2022
360. Title: Robust Multi-Robot Active Target Tracking Against Sensing and Communication Attacks, Venue: IEEE ACC 2022, Lead Author: Lifeng Zhou, Year: 2022
361. Title: Sequencing of multi-robot behaviors using reinforcement learning, Venue: Control Theory and Technology, Lead Author: P. Pierpaoli, Year: 2021
362. Title: Stability of Neural Networks on Manifolds to Relative Perturbations, Venue: ICASSP 2022, Lead Author: Zhiyang Wang, Year: 2022
363. Title: Stable and Transferable Wireless Resource Allocation Policies via Manifold Neural Networks, Venue: ICASSP 2022, Lead Author: Zhiyang Wang, Year: 2022
364. Title: Training Stable Graph Neural Networks Through Constrained Learning, Venue: ICASSP 2022, Lead Author: Juan Cervino, Year: 2022

365. Title: The Situation Awareness Framework for Explainable AI (SAFE-AI) and Human Factors Considerations for XAI Systems, Venue: International Journal of Human-Computer Interaction, Lead Author: Lindsay Sanneman, Year: 2022

List of Symbols, Abbreviations, and Acronyms

1-D	1-dimensional
3-D	3-dimensional
AI	artificial intelligence
AIA	Active Information Acquisition
AIMM	ARL's AI for Maneuver and Mobility
ARL	Army Research Laboratory
BADGR	Berkeley Autonomous Driving Ground Robot
BPP	Biennial Program Plan
CDE	Cross-Disciplinary Research Experiment
CLEAR	Consistent Lifting, Embedding, and Alignment Rectification
CLIPPER	Consistent Linking, Pruning, and Pairwise Error Rectification
CPU	central processing unit
CRA	Collaborative Research Alliance
DEVCOM	US Army Combat Capabilities Development Command
DCIST	Distributed and Collaborative Intelligent Systems and Technology
EOT	Emerging Overmatch Technologies
ERP	Essential Research Program
FY	fiscal year
GIS	Geographic Information System
GNN	graph neural network
GPS	global positioning system
GrBAL	gradient-based adaptive learner
HRT	human–robot teaming
ISR	intelligence, surveillance, and reconnaissance
K-NODE	knowledge-based neural ordinary differential equation

LaND	Learning to Navigate from Disengagements
LDPC	low-density parity checks
LSD	Learned Sampling Distribution
LTL	Linear Temporal Logic
MADDPG	multi-agent deep deterministic policy gradient
MARL	multi-agent reinforcement learning
MDO	Multi-Domain Operations
MDS	maximum distance separable
MID	Mobile Infrastructure on Demand
MIT	Massachusetts Institute of Technology
ML	machine learning
MPPI	model predictive path integral control
MSF	mean safe flight
MTRL	multi-task reinforcement learning
OrcVIO	Object residual constrained Visual Inertial. Odometry
PAC	Perception-Action-Communication
PGO	pose graph optimization
RAS	Robotics and Autonomous Systems
ReBAL	recurrence-based adaptive learner
RECON	Rapid Exploration for Open-World Navigation
RL	reinforcement learning
SLAM	simultaneous localization and mapping
STRATA	Stochastic Trait-based Task Assignment
TPR	Technical Progress Report
TMG	Technical Management Group
UAV	Unmanned Aerial Vehicle
VICTOR	Versatile Tactical Power and Propulsion

1 (PDF)	DEFENSE TECHNICAL INFORMATION CTR DTIC OCA	FCDD RLD E KS FOSTER FCDD RLD F K KAPPRA FCDD RLD FR M TSCHOPP FCDD RLD SM L BLUM FCDD RLH J CHEN PJ FRANASZCZUK C LANE K MCDOWELL FCDD RLH B JJ SUMNER FCDD RLH F JR GASTON FCDD RLH T D STRATIS-CULLUM FCDD RLL T KINES FCDD RLL D J S ADAMS FCDD RLL DP J MCCLURE FCDD RLR B HALPERN S LEE D STEPP FCDD RLR E RA MANTZ C VARANASI FCDD RLR EL JX QIU MD ULRICH FCDD RLR EN RA ANTHENIEN JR FCDD RLR IC MA FIELDS SP IYER FCDD RLR IM JD MYERS FCDD RLR IN XN WANG FCDD RLR P REYNOLDS FCDD RLR P LL TROYER FCDD RLR PC D POREE FCDD RLR PL MK STRAND FCDD RLS J ALEXANDER M GOVONI M WRABACK
1 (PDF)	DEVCOM ARL FCDD RLD DCI TECH LIB	
1 (PDF)	DA HQ DASA(R&T)	
9 (PDF)	USARMY AFC L BROUSSEAU J REGO A LINZ K WADE S BRADY J REGO T KELLY E JOSEPH B SESSLER	
2 (PDF)	DEVCOM HQ FCDD ST C SAMMS M HUBBARD	
77 (PDF)	DEVCOM ARL FCDD RLC C BEDELL B SADLER B PIEKARSKI H EVERITT FCDD RLC CA L KAPLAN FCDD RLC ES G VIDEEN S HILL Y PAN FCDD RLC I FCDD RLC IA J FINK E STUMP FCDD RLC IS M DOROTHY B MACCALL FCDD RLC N BM RIVERA A SWAMI FCDD RLD P BAKER A KOTT S SILTON FCDD RLD D T ROSENBERGER	

FCDD RLS C
M REED
FCDD RLS CC
S BEDAIR
FCDD RLS CE
TR JOW
K XU
FCDD RLS CL
M DUBINSKIY
FCDD RLS E
RD DELROSARIO
FCDD RLS ED
K JONES
FCDD RLS EA
A ZAGHLOUL
FCDD RLS S
WL BENARD
FCDD RLS SO
W ZHOU
FCDD RLW
S KARNA
JF NEWILL
AM RAWLETT
SE SCHOENFELD
J ZABINSKI
FCDD RLW B
R BECKER
FCDD RLW C
C KRONINGER
FCDD RLW M
ES CHIN
FCDD RLW S
V CHAMPAGNE
AL WEST
FCDD RLW T
RZ FRANCART
FCDD RLW TC
JD CLAYTON
FCDD RLW W
TV SHEPPARD
FCDD RLW WA
B RICE
R PESCE-RODRIGUEZ
FCDD RLW M
A HALL