ARL-TR-10274 • JAN 2026

The Army Artificial Intelligence Innovation Institute (Summary Technical Report, 2018–2025)

by Mark A. Tschopp and Daniel N. Cassenti

NOTICES

Disclaimers

ARL-TR-10274 ● JAN 2026

The Army Artificial Intelligence Innovation Institute (Summary Technical Report, 2018–2025)

by Mark A. Tschopp and Daniel N. Cassenti
DEVCOM Army Research Laboratory

REPORT DOCUMENTATION PAGE

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED	
		START DATE	END DATE
January 2026	Summary Technical Report	01/01/2018	12/31/2025

4. TITLE AND SUBTITLE
The Army Artificial Intelligence Innovation Institute (Summary Technical Report, 2018–2025)

5a. CONTRACT NUMBER	5b. GRANT NUMBER	5c. PROGRAM ELEMENT NUMBER
5d. PROJECT NUMBER	**5e. TASK NUMBER**	**5f. WORK UNIT NUMBER**

6. AUTHOR(S)
Mark A. Tschopp and Daniel N. Cassenti

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)	8. PERFORMING ORGANIZATION REPORT NUMBER
DEVCOM Army Research Laboratory ATTN: FCDD-RLA-I 2800 Powder Mill Rd, Adelphi, MD 20783	ARL-TR-10274

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
DISTRIBUTION STATEMENT A. Approved for public release: distribution unlimited.

13. SUPPLEMENTARY NOTES
ORCID IDs: Mark A. Tschopp, 0000-0001-8471-5035; Daniel N. Cassenti, 0000-0002-8139-5993;
email: mark.a.tschopp.civ@army.mil.

14. ABSTRACT
The Army Artificial Intelligence Innovation Institute (A2I2) collaborative research alliance was established in FY19 to advance artificial intelligence (AI) and machine learning (ML) for Army modernization. Focused on foundational techniques such as computer vision, autonomy, and adversarial robustness, the program evolved to incorporate large language models (LLMs) for language-grounded planning, open-vocabulary perception, and human–AI interaction. A2I2 explored robust decision-making in contested environments, self-supervised learning, multiagent systems, and integration of LLMs and vision–language models into pipelines for perception, planning, and decision support. Key outcomes include advances in autonomous systems for off-road navigation, task planning, and manipulation; multimodal perception and information fusion; and reusable datasets and benchmarks to accelerate model development and evaluation. The program produced over 300 publications, over 4,000 citations, and impactful open-source software and tools, fostering collaboration across academia, industry, and Army researchers. A2I2 will continue to address Army-relevant challenges related to AI in the physical, information, and cognitive domains, delivering capabilities that enhance decision-making, situational awareness, and operational effectiveness across Army missions.

15. SUBJECT TERMS
artificial intelligence; machine learning; autonomy; robotics; large language models; information fusion; vision-language models; decision support; computer vision; perception; situational awareness; Military Information Sciences; Science of Intelligent Systems; Summary Technical Report

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES
a. REPORT	b. ABSTRACT	C. THIS PAGE		
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UU	120

19a. NAME OF RESPONSIBLE PERSON	19b. PHONE NUMBER *(Include area code)*
Mark A Tschopp	301-938-5389

STANDARD FORM 298 (REV. 5/2020)
Prescribed by ANSI Std. Z39.18

Contents

List of Figures

List of Tables

Executive Summary

Artificial intelligence (AI) and machine learning support the Army's decision advantage, resilient autonomy, and mission effectiveness in Multi-Domain Operations (MDO). The Army Artificial Intelligence Innovation Institute (A2I2), formulated in 2018 and established in FY19 as a collaborative research alliance under U.S. Army Combat Capabilities Development Command (DEVCOM) Army Research Laboratory's (ARL's) Broad Agency Announcement portfolio, initially focused on two operationally critical challenge areas: Expeditionary Maneuver and Air/Ground Reconnaissance. Early efforts prioritized perception, planning and control, robustness, and data-efficient learning on edge compute; then adapted midstream as large language models (LLMs) matured, integrating language-grounded planning, open-vocabulary semantics, and human–AI teaming to support embodied AI at the tactical edge.

A2I2 advanced foundational and applied research by linking performers with Army problem owners, field-relevant data, and representative evaluation environments. Research emphasized autonomy in GPS-degraded and communications-constrained conditions; learning physics-informed abstractions; depth completion and traversability estimation; task and motion planning among movable objects; adversarial robustness with provable methods; personalized federated learning; and privacy-preserving analytics with unlearning verification. A2I2 coordinated 24 active awards across academia and industry, in partnership with ARL scientists, moving concepts toward capability via open-source software, reusable datasets, and rigorous evaluation.

Key outcomes underscore breadth, quality, and Army relevance:

- Program scale and collaboration: 311 publications authored by 742 unique contributors across 116 institutions; 65 ARL researchers coauthored 110 publications; 89% of papers are multi-institution and 89% are U.S.-only, reflecting strong domestic collaboration with targeted international engagement.
- Scientific influence: 4,266 cumulative citations and an h-index of 30 (i.e., 30 papers with at least 30 citations), with 22 papers in the top 1% and 109 in the top 10% by topic area and venue, indicating sustained impact and community uptake.
- Transition enablers: Open-source releases, benchmarks (e.g., standardized adversarial robustness), large-scale datasets (e.g., off-road dynamics), and GitHub

repositories that support reproducibility and comparative evaluation across representative scenarios.

Operationally relevant advances span core Army missions (parenthetical references point to top-cited articles and GitHub repositories):

- Expeditionary Maneuver: High-speed off-road autonomy via inverse kinodynamics; predictive traversability and self-supervised costmaps; unsupervised depth completion with calibrated backprojection; adaptive planner parameter learning for dynamic environments; planning and manipulation among movable objects for logistics and preparation tasks.
- Air/Ground Reconnaissance: Open-vocabulary 3D scene graphs (ConceptGraphs) for semantic perception and spatial reasoning; multi-view action recognition via contrastive learning; lightweight, real-time tracking on edge devices; adversarially robust detection (DetectorGuard, PatchGuard) with provable methods; standardized robustness evaluation (RobustBench) to inform model selection for deployment.
- Distributed learning, privacy, and cyber: Hierarchical, personalized federated learning orchestrated via UAVs in multi-tier networks; systematic evaluation of membership inference and per-sample privacy risk; probabilistic verification of machine unlearning (Athena); automated cyber threat hunting that operationalizes open-source threat intelligence.
- Human–AI teaming with LLMs: Language-grounded tasking and re-planning (ProgPrompt, LLM-Planner) to improve adaptability and reduce operator burden in contested, partially observable environments.

Dissemination and visibility reflect a balance between top-tier technical venues (e.g., ICRA, CVPR, ICCV, IROS) and domain-focused forums (e.g., HPEC, SPIE AI/ML for MDO), supporting scientific rigor and DoD engagement. Most outputs carry DOIs, supporting discoverability and citation tracking, and many publications acknowledge A2I2 support, documenting Army investment in open science.

Looking ahead, A2I2's approach (building on solid technical foundations while integrating new advances) supports Army efforts to develop assured, data-efficient, and trustworthy autonomy. Priority directions include privacy-preserving edge AI; LLM-enabled perception, planning, and mission management; scalable human–machine teaming; and evaluation under adversarial, GPS-degraded, and communications-constrained conditions. By sustaining open science, operationally grounded research, and agile program execution, A2I2 will continue to advance and transition AI systems to support decision-making across Army missions and domains.

1. Introduction: Why AI for the Army

Artificial intelligence (AI) and machine learning (ML) support faster decision-making, improved situational awareness, and mission effectiveness across Multi-Domain Operations (MDO). As battlespaces grow more complex and contested—from electromagnetic warfare to GPS-denied, unstructured terrain—AI enables autonomy and decision support at the tactical edge. Systems that can perceive, reason, learn, and adapt under uncertainty extend Soldier reach, reduce operator burden, and contribute to outcomes in ISR (Intelligence, Surveillance, and Reconnaissance), logistics, maneuver, C2 (Command and Control), and cyber defense. AI also strengthens human–AI teaming by enabling autonomous systems to collaborate with Soldiers under operational constraints.

The 2010s marked a foundational period for modern AI. Convolutional neural networks (CNNs) substantially improved object recognition, detection, and segmentation, enabling practical perception for autonomous systems. In parallel, representation learning mapped unstructured text into embeddings (dense numerical vectors that capture semantic relationships), supporting search, retrieval, reasoning, and grounding between language and perception. Together, these advances improved environmental understanding for robotics and ISR: systems could detect and track objects, estimate depth and traversability, and fuse multi-sensor data more effectively. These capabilities emerged from both disruptive breakthroughs (e.g., deep architectures, large-scale training) and steady engineering improvements (e.g., optimization, data curation, rigorous evaluation), a trajectory expected to change how the Army fights and equips Soldiers.

Within this evolving landscape, the Army Artificial Intelligence Innovation Institute (A2I2) was formulated in 2018 and established in FY19 under U.S. Army Combat Capabilities Development Command (DEVCOM) Army Research Laboratory's (ARL's) Broad Agency Announcement (BAA) portfolio as a collaborative research alliance (CRA) to advance AI/ML for two challenge areas: *Expeditionary Maneuver* and *Air/Ground Reconnaissance*. For Expeditionary Maneuver, the program emphasized systems that can reason about, interact with, and navigate complex environments, pursuing research in: learning physics-based abstractions, modeling human and object dynamics, interactive scene understanding, and maneuver planning under uncertainty. For Air/Ground Reconnaissance, the program studied how diverse, embodied agents can collectively sense, infer, reason, and plan with Soldiers, focusing on: self-organizing teams, adaptive systems, multi-level human–agent interaction, and optimizing perception–action–communication loops.

Because the program launched prior to the rise of today's large language models (LLMs), early priorities centered on perception, autonomy, robustness, and data-efficient learning on edge compute. As LLMs matured and proved useful for planning, reasoning, and open-vocabulary perception, the program adapted: investigators integrated language-grounded tasking, semantic scene understanding, and human–robot teaming into embodied AI pipelines. This reflects A2I2's core approach—build on solid technical foundations while incorporating emerging advances that improve operational utility.

These research advances require collaboration. The A2I2 operates through cooperative agreements and partnerships that connect Army problem owners with academic and industry teams, combining domain insight, scientific rigor, and engineering capability. This ensures research is driven by operational needs, evaluated on representative data and tasks, and moved from concept toward capability through open-source code, reusable datasets, and comparative evaluation. This report summarizes Phase I of A2I2, situating scientific context from the 2010s through program formulation in 2018 and outlining why Army investment in AI aims to provide data-efficient, operationally relevant capabilities across MDO.

2. Technical Summary

2.1 Motivation and History

The A2I2 CRA was established in FY19 to address gaps in AI capabilities for future Army operations. The program responded to the need to advance AI for Autonomous Maneuver in support of MDO. In the 2010s, deep learning transformed perception and representation: CNNs improved visual understanding; representation learning turned unstructured text and multimodal signals into embeddings that could be searched, fused, and reasoned over; and robust, data-efficient methods began to enable autonomy in unstructured environments. Formulated in 2018, A2I2 emphasized perception, planning and control, robustness against adversaries, and data-efficient learning on edge compute. As LLMs matured, the program integrated language-grounded planning and open-vocabulary perception into embodied systems, preserving operational focus while leveraging emerging advances.

The original motivation stems from two fundamental challenges:

- **Expeditionary Maneuver:** Need for systems that can reason about, interact with, navigate through, and manipulate dynamic environments to achieve Army-relevant actions in future multi-domain battlespace.
  - Operate in GPS-degraded, communications-constrained, and adversarial electromagnetic conditions with limited compute, power, and bandwidth.
  - Perceive and fuse multi-sensor inputs (RGB-D, LiDAR, EO/IR, RF) to infer traversability, geometry, semantics, and affordances in open-world, off-road, and urban terrains.
  - Plan and control at speed under uncertainty, adapting online to distribution shift, environmental change, and mission updates.
- **Air/Ground Reconnaissance:** Requirement for diverse, embodied agents that can collectively sense, infer, reason, plan, and execute with Soldiers while facing peer adversaries.
  - Build ISR pipelines for detection, tracking, and activity recognition across viewpoints and modalities under label scarcity and contested conditions.
  - Ensure robustness, privacy, and security (adversarial resilience, federated learning integrity, privacy-preserving analytics) to protect mission data and models.
  - Enable human–AI teaming through interpretable behaviors and language-grounded tasking as LLMs and vision–language models matured.

2.2 Vision and Goals

The A2I2 vision is to develop and employ a suite of AI to provide intelligent systems that assist Soldiers in dynamic, uncertain, complex, contested, austere, and congested operational environments. The intelligent systems will reduce the dangerous, physically demanding, and routine tasks required of Soldiers and crew members. The A2I2 will develop intelligent agents for the battlefield (heterogeneous and distributed) that rapidly learn, adapt, reason, and act in these environments. Specifically, AI agents that 1) learn, adapt, reason, and act using small datasets of noisy, dynamic, deceptive data; 2) learn, adapt, and reason in a dynamic, distributed manner over highly heterogeneous data with impaired command, control, communications, computers, and intelligence (C4I); 3) perform computations under extreme size, weight, and power, time (SWaPT) constraints; and 4) process at the point of need.

AI will provide our Warfighters with the ability to recognize, with one glance, the tactical advantages and disadvantages on the battlefield using a heterogeneous mix of unmanned ground and aerial platforms that rapidly learn, adapt, and reason in complex and dynamic environments. These physical agents will interact with Soldiers as teammates with a shared understanding to jointly respond to dynamic events while encountering adversarial deception at the speed of the fight. Autonomous unmanned systems will respond to digital and verbal commands and act as members of the squad or crew. They will provide accurate verbal and written language translation unobtrusively. Autonomous unmanned systems will function as members of the formation, executing tasks as well as providing oversight for subordinate systems. This capability will allow leaders to employ unmanned systems for critical and complex tasks such as establishing a mesh communication network or reconnoitering and mapping subterranean infrastructures.

A2I2 will focus on addressing unique Army AI technical challenges. The overarching A2I2 goals are to 1) focus AI researchers on technical challenges faced by future Army MDO; 2) serve as a focal point and catalyst for Army AI innovation; 3) develop a fundamental understanding of AI in order to predict model performance in various MDO phases, domains, geographies, and environments; 4) create a robust training, testing, and experimentation process for rapid assessment in Army-relevant MDO conditions; and 5) rapidly mature promising research algorithms and models for transition to the Warfighter.

2.3 Strategy

A2I2 pursues an integrated strategy for accelerating the fundamental understanding, development, training, testing, experimentation, and transition of AI for Army MDO. ARL uses collaborative alliances to address research problems that require a multidisciplinary research approach to make significant progress. These alliances are uniquely designed to 1) accelerate the development of AI/ML for the Army with a model that enables government–private-sector collaboration and exchange, 2) create an open research environment where ideas are shared freely, and 3) flexibly adjust research efforts on the fly to adapt to emerging trends.

Alliances consist of a competitively selected consortium of academia and industry that, when combined with ARL, forms an alliance. The alliance is structured and managed to foster an open collaborative research environment to allow deep collaboration among government and private-sector researchers to conduct multidisciplinary research. This is enabled by an innovative acquisition model that allows and encourages government scientists to collaborate, co-invent, and co-publish with private-sector researchers. To develop a shared understanding and to cultivate these collaborations, the alliance encourages frequent technical interchanges and long- and short-term staff rotations.

A series of competitively selected basic and applied research tasks jointly developed by the alliance encourages new ideas, allows new researchers to join the alliance, and adapts to emerging technical gaps and scientific breakthroughs. The innovative use of agreements affords ARL significant flexibilities to 1) adjust alliance research on the fly without contractual negotiations, 2) add and remove research organizations and researchers to accommodate new innovative ideas or emerging research trends, 3) augment research with government staff as they become available, and 4) fully embed alliance research in the ARL enterprise. This acquisition agility is critically important for impacting the rapidly expanding field of AI/ML.

Initial Program Focus (FY19–FY25)

The initial program focus of the A2I2 was:

1. **Perception and Representation at the Edge:** Data-efficient self-/unsupervised methods for depth completion, traversability estimation, multi-view action recognition, and open-world semantics; creation and

curation of Army-relevant datasets and scene representations to accelerate embodied AI.

2. **Planning, Control, and Manipulation:** Task and motion planning among movable objects; adaptive planner parameter learning; high-speed off-road kinodynamics with real-time feedback and safety constraints.

3. **Robustness, Security, and Assurance:** Standardized robustness benchmarks and provable methods (e.g., adversarial patches, detector hardening); quantification of privacy risk and unlearning verification for distributed and sensitive data regimes.

4. **Distributed and Resource-Constrained Learning:** Hierarchical, personalized federated learning and online adaptation across multi-tiered networks and UAV/UGV platforms, mindful of communication, power, and compute limits.

5. **Transition Enablers:** Open-source code, reusable datasets, and comparative evaluation across representative scenarios to shorten the path from concept to practice; maintained repositories that support reproducibility and re-use.

A2I2 relied on multidisciplinary collaboration among academia, industry, and Army researchers, tightly coupled to operational problems. This kept scientific innovation grounded in realistic constraints and produced algorithms, software prototypes, benchmarks, and datasets with clear paths to evaluation and transition. The program's motivation and early focus reflect a deliberate strategy: invest in foundational perception, planning, robustness, and data-efficient learning for expeditionary autonomy; adapt as the AI landscape evolves (e.g., integrating LLMs for planning and semantics); and provide AI capabilities that support decision-making and maneuver across Army missions.

2.4 Technical Efforts and Methods

As noted earlier, A2I2 conducts research via cooperative agreements (CAs) that permit government participation in planning and execution. A biannual planning cycle enabled adjustments to emerging needs and technology shifts. In the initial phase, 24 active awards across academia and industry were utilized to team with ARL scientists for moving ideas toward application: algorithms, software, benchmarks, and datasets.

Technical Focus Areas: Initial emphases centered on expeditionary maneuver and air/ground reconnaissance, reflecting applied research challenges in the circa-2018 time frame (i.e., prior to the large language model disruption). Representative early focus areas included:

Expeditionary Maneuver Research:

- Learning physics abstractions from observations to support prediction and control under uncertainty.
- Modeling complex objects and human interactions for manipulation and scene-aware planning.
- Interactive perception to improve physical scene understanding.
- Planning and control in GPS-degraded and communications-constrained conditions.

Air/Ground Reconnaissance Research:

- Adaptive systems that sustain perception and inference under distribution shift.
- Team coordination based on mission goals, platform capabilities, and communications.
- Multi-level human–AI interaction for tasking, explanation, and shared situational awareness.
- Optimization of perception–action–communication loops at the tactical edge.

2.5 Key Findings and Metrics

The A2I2 has demonstrated significant impact in advancing AI research and fostering collaboration across institutions. Since inception, the program produced 311 publications authored by 742 unique contributors from 116 institutions worldwide. Sixty-five ARL researchers contributed to 110 of these publications, underscoring a strong intramural presence. Collaboration has been a cornerstone of the program: 89% of papers involve multiple institutions, and 89% are authored exclusively by U.S.-based organizations. The program's bibliometric achievements evidence quality and influence: 22 papers rank in the top 1% of citations (by topic area and venue), and 109 papers rank in the top 10%. Cumulative citations total 4,266, with an h-index of 30 (i.e., 30 papers with at least 30 citations), reflecting durable impact and community uptake.

- Number of publications: 311
- Number of unique authors: 742
- Number of unique ARL authors: 65 (on 110 publications)
- Number of unique institutions: 116
- Percentage of multi-institution papers: 89%
- Percentage of papers w/ only U.S. institutions: 89%
- Number of papers in top 1%: 22
- Number of papers in top 10%: 109
- Number of citations: 4,266
- h-index: 30 (30 papers with at least 30 citations)

Collectively, these outcomes reflect the program's success in producing high-impact research that advances Army objectives in AI/ML and accelerates transition through open-source code, datasets, and rigorous evaluation (see Appendix A).

2.6 Scientific and Technical Impact

This section presents synopses of highly cited research papers that contribute to the A2I2 h-index, showcasing the breadth and depth of scientific contributions in AI and autonomous systems.* Each synopsis follows a standardized format that includes the paper title, a focused research question derived from the work, and direct links to both the Digital Object Identifier (DOI) for accessing the full publication and associated GitHub repositories where available. The summaries provide concise technical overviews of the methodologies and key findings, followed by an assessment of the work's relevance to Army operations and mission requirements. When applicable, transition pathways are noted, often facilitated through open-source code repositories and datasets made available through the linked GitHub

*During the preparation of this work, the authors used AskSage to assist in summarizing the findings within these articles. Please use the DOI link to access the full text of the articles.

sites, enabling broader adoption and further development of these research contributions.

The papers span critical areas in artificial intelligence including autonomous navigation, adversarial robustness, federated learning, computer vision, and cybersecurity, demonstrating the multidisciplinary nature of modern AI research and its applications to defense challenges. Each entry represents work that has achieved significant citation impact, indicating its influence on the broader research community and potential for operational transition.

ProgPrompt: Generating Situated Robot Task Plans with LLMs

Research Question: How can large language models (LLMs) be leveraged to generate effective and context-aware task plans for robots in situated environments?

DOI Link GitHub Repository The ProgPrompt paper explores the use of large language models (LLMs) for generating task plans for robots operating in situated environments. The authors propose a framework called ProgPrompt, which integrates LLMs with domain-specific programming constructs to produce executable plans tailored to robotic tasks. The methodology involves prompting LLMs with structured templates and contextual information, enabling them to generate plans that are both syntactically correct and semantically meaningful. Experimental results demonstrate the effectiveness of ProgPrompt in diverse scenarios, including navigation and object manipulation, highlighting its adaptability and robustness. The study also addresses challenges such as grounding abstract instructions into actionable steps and ensuring compatibility with robotic control systems.

For the Army, this framework enables context-aware task planning for autonomous ground and air systems conducting reconnaissance, logistics, and search-and-rescue in contested, GPS-denied environments. It improves mission tempo and reliability while reducing operator burden, and is positioned for transition through open-source code and modular integration with existing autonomy stacks.

LLM-Planner: Few-Shot Grounded Planning for Embodied Agents with Large Language Models

Research Question: How can large language models be used for few-shot planning in embodied agents that follow natural language instructions in complex, partially observable environments?

DOI Link GitHub Repository The LLM-Planner paper proposes a grounded re-planning algorithm that dynamically updates high-level plans during execution by incorporating environmental observations (detected objects) into LLM prompts. This innovation enables the system to adapt plans based on what the agent perceives in real-time, rather than relying on static, pre-generated plans. The approach achieves competitive performance using less than 0.5% of training data compared to fully trained baselines on the ALFRED dataset.

For the Army, this technology facilitates rapid deployment of autonomous systems capable of understanding natural language commands and adapting their behavior dynamically, making them ideal for battlefield scenarios where pre-programmed responses may be insufficient and quick adaptation to new situations is critical.

UHP-SOT++: An Unsupervised Lightweight Single Object Tracker

Research Question: How can unsupervised, lightweight single object trackers achieve high performance while maintaining real-time operation on resource-limited platforms?

DOI Link The UHP-SOT++ paper proposes an enhanced fusion strategy that systematically combines proposals from three complementary modules: discriminative correlation filters for appearance modeling, background motion modeling for global scene understanding, and trajectory-based box prediction for maintaining object motion inertia. This approach enables the tracker to recover from tracking loss and adapt to rapid object deformations without requiring pre-training on large datasets.

For the Army, this technology provides an ideal solution for real-time object tracking on drones, mobile platforms, and edge devices where computational resources are limited, enabling autonomous surveillance, target tracking, and reconnaissance missions while operating at 20 FPS on standard CPU hardware without requiring GPU acceleration or extensive training data.

RobustBench: A standardized adversarial robustness benchmark

Research Question: How can adversarial robustness be evaluated in a standardized, reliable manner to accurately track progress and identify the most promising approaches in training robust models?

DOI Link GitHub Repository The RobustBench paper establishes a standardized benchmark that addresses robustness evaluation challenges through three key innovations: 1) clearly defined restrictions on model types to prevent gradient obfuscation and ensure reliable evaluation, 2) AutoAttack as a parameter-free ensemble of diverse attacks for consistent robustness assessment, and 3) a comprehensive Model Zoo providing unified access to 80+ robust models. This approach enables systematic comparison of adversarial defenses while preventing overestimation of robustness that plagued previous evaluations.

For the Army, this benchmark provides a critical tool for evaluating and selecting robust AI systems for deployment in adversarial environments, ensuring that defensive capabilities are accurately assessed before integration into mission-critical applications where model reliability under attack is paramount for operational success.

ConceptGraphs: Open-Vocabulary 3D Scene Graphs for Perception and Planning

Research Question: How can robots build scalable, semantically rich 3D representations of their environment that support open-vocabulary perception, spatial reasoning, and language-conditioned task planning?

DOI Link GitHub Repository ConceptGraphs introduces an open-vocabulary, object-centric 3D scene graph representation that fuses geometric information from RGB-D data with semantic cues from large vision–language models. Instead of storing dense per-point features, ConceptGraphs segments objects in 2D, associates and fuses them across viewpoints, and generates a 3D graph where each node carries captions, spatial attributes, and affordance-relevant semantics. Object relationships are inferred using large language models, enabling robots to reason about spatial structure ("A is on B") and utility (e.g., traversability, helpfulness, affordances). The resulting scene graph supports complex natural-language queries—such as finding objects by description, identifying functional substitutes, or determining where a robot may navigate or manipulate—and enables downstream tasks including segmentation, grounding, navigation, and manipulation across wheeled and legged robot platforms.

For the Army, this technology provides a foundation for autonomous perception and planning in complex, dynamic, and unstructured environments. By enabling robots to understand scenes through open-vocabulary semantics and relational structure, ConceptGraphs enhances mission-critical capabilities such as reconnaissance, route planning, object search and retrieval, and human–robot teaming. This supports more adaptive robotic behaviors in multi-domain operations without requiring extensive task-specific training data.

Enabling Efficient Cyber Threat Hunting With Cyber Threat Intelligence

Research Question: How can open-source Cyber Threat Intelligence (OSCTI) be leveraged to make cyber threat hunting in system audit logs more efficient, accurate, and automated?

DOI Link GitHub Repository The *ThreatRaptor* system bridges external threat intelligence with internal system logs. It builds a lightweight, unsupervised NLP pipeline to extract structured threat behaviors (e.g., Indicators of Compromise and their interrelations) from unstructured OSCTI text. ThreatRaptor introduces a domain-specific query language called TBQL (Threat Behavior Query Language), which is expressive yet concise, enabling users to define behavioral threat patterns. The system then uses a query-synthesis mechanism to automatically generate TBQL queries from the extracted threat behavior graphs. Finally, it executes these queries efficiently over large audit-log datasets via a high-performance engine. Experimental evaluation shows that ThreatRaptor achieves high accuracy and efficiency in realistic attack scenarios, enabling threat hunters to identify malicious system activities without extensive manual query writing. The authors report strong performance on extraction (e.g., high F1) and robust threat hunting across a broad set of attack cases.

For the Army, this capability offers significant operational value: by automating the translation of threat intelligence into actionable log queries, ThreatRaptor can speed up threat hunting, reduce reliance on human-crafted hunts, and adapt quickly as new threat intelligence emerges. Such a system strengthens continuous defense by tightly coupling strategic threat intel with tactical detection capabilities.

DefakeHop: A Light-Weight High-Performance Deepfake Detector

Research Question: How can we detect deepfake face manipulations accurately and efficiently using a model that is small enough to run on resource-constrained devices?

DOI Link GitHub Repository DefakeHop proposes a novel, lightweight deepfake detection algorithm based not on large neural networks, but on the successive subspace learning (SSL) principle. The method uses c/w-Saab transforms to extract unsupervised features from different facial regions, then applies a feature distillation module that reduces spatial dimensionality and does a "soft" classification per channel to create compact but discriminative face descriptors. These features are then classified to distinguish real versus manipulated content. Despite its very small model size (only 42,845 parameters), DefakeHop achieves strong performance: on benchmark datasets, it reports AUCs of 100% on UADFV, 94.95% on Celeb-DF v1, and 90.56% on Celeb-DF v2.

For the Army, this technology is especially valuable: it provides a low-latency, low-compute deepfake detection capability that could be embedded onboard edge devices (e.g., Soldier-worn cameras or tactical systems). This enables real-time verification of visual intelligence and guarding against adversarial misinformation or spoofed imagery under constrained bandwidth and compute conditions.

Systematic Evaluation of Privacy Risks of Machine Learning Models

Research Question: How can we systematically measure and benchmark the privacy risks of machine learning models, especially regarding membership inference attacks, in a fine-grained manner beyond aggregated metrics?

DOI Link GitHub Repository Song and Mittal propose a rigorous evaluation framework for membership inference attacks. They argue that prior work underestimates privacy risk by only training neural-network attack models and focusing on aggregate metrics like overall attack accuracy. Instead, they develop a suite of benchmark attacks—including non-neural methods and a novel entropy-based attack—to more faithfully assess model vulnerability. They also define a new metric, the privacy risk score, which estimates for each individual sample the likelihood it was part of the training set, allowing an adversary to target high-risk data points with confidence. Their experimental analysis shows that existing defenses against membership inference (e.g., regularization, differential privacy mechanisms) are less effective than previously thought, especially when evaluated against adaptive adversaries. They also find strong heterogeneity in risk: some samples are far more vulnerable than others, and high-risk samples correlate with model sensitivity, generalization error, and feature embeddings.

For the Army, this work has strong relevance: it reveals how models trained on sensitive or classified data (e.g., personnel records, mission logs, ISR imagery) may leak membership information. By applying such a systematic risk assessment, Army AI programs can identify which data points or subpopulations are most at risk and prioritize the use of stronger defenses. This supports more robust, privacy-aware model deployment in operational and training environments.

Feature 3DGS: Supercharging 3D Gaussian Splatting to Enable Distilled Feature Fields

Research Question: How can 3D Gaussian Splatting be enhanced to efficiently represent and distill feature fields for 3D scene reconstruction?

DOI Link GitHub Repository The Feature 3DGS paper introduces an advanced method for 3D Gaussian Splatting that incorporates distilled feature fields to improve the representation and reconstruction of 3D scenes. The approach leverages neural networks to optimize splatting parameters and distill high-dimensional features into compact representations. The authors demonstrate significant improvements in rendering quality and computational efficiency compared to existing methods. Experimental results show that Feature 3DGS achieves state-of-the-art performance on benchmark datasets, with reduced memory overhead and faster processing times. The method is particularly effective in handling complex scenes with intricate geometric and photometric details.

For the Army, this capability strengthens 3D scene reconstruction for autonomous navigation, surveillance, and situational awareness in complex and GPS-denied environments. It delivers more efficient and accurate 3D modeling to support ISR, enabling faster, more reliable decision-making at the tactical edge.

Unsupervised Depth Completion with Calibrated Backprojection Layers

Research Question: How can unsupervised learning methods improve depth completion accuracy using calibrated backprojection layers?

DOI Link GitHub Repository This paper presents a novel approach to depth completion using unsupervised learning techniques combined with calibrated backprojection layers. The method leverages geometric constraints and self-supervised signals to refine depth predictions from sparse input data. By introducing calibrated backprojection layers, the authors ensure accurate alignment between depth maps and image features, improving reconstruction quality. Experimental results demonstrate that the proposed method achieves competitive performance on benchmark datasets without requiring ground truth depth labels, making it suitable for scenarios with limited labeled data. The approach is computationally efficient and robust to noise in sparse depth inputs.

For the Army, this work enhances depth sensing for autonomous navigation, target detection, and situational awareness in complex, contested, and GPS-denied environments. Its unsupervised methods reduce dependence on labeled data, enabling rapid deployment across diverse operational scenarios and edge platforms where data collection is limited.

Planning for Manipulation among Movable Objects: Deciding Which Objects Go Where, in What Order, and How

Research Question: How can planning algorithms optimize the manipulation of movable objects by determining their placement, sequence, and manipulation strategies?

DOI Link GitHub Repository This paper addresses the problem of planning for manipulation tasks involving movable objects, focusing on determining optimal object placements, sequencing, and manipulation strategies. The authors propose a novel algorithmic framework that integrates task and motion planning to handle complex environments with multiple movable objects. The approach leverages heuristic search methods and geometric reasoning to efficiently compute feasible plans. Experimental evaluations demonstrate the framework's ability to solve challenging manipulation scenarios with improved computational efficiency and success rates compared to existing methods. The work contributes to advancing autonomous robotic manipulation capabilities in dynamic and cluttered environments.

For the Army, this research advances robotic manipulation for logistics, equipment handling, and operational support in constrained, dynamic, and contested environments. Improved planning algorithms enable robots to sequence and execute complex tasks—supply movement, battlefield preparation, and handling of hazardous or bulky materiel—with greater efficiency, reliability, and reduced Soldier workload.

Robust and Accurate Object Detection via Adversarial Learning

Research Question: How can adversarial learning improve the robustness and accuracy of object detection models?

DOI Link GitHub Repository This paper explores the use of adversarial learning techniques to enhance the robustness and accuracy of object detection models. The authors propose a framework that integrates adversarial training with feature refinement to mitigate the impact of adversarial perturbations on detection performance. By generating adversarial examples during training, the model learns to adapt to challenging scenarios and improves its generalization capabilities. Experimental results on standard object detection benchmarks demonstrate significant improvements in robustness against adversarial attacks while maintaining high detection accuracy. The approach is computationally efficient and scalable to large datasets.

For the Army, this technology hardens object detection for autonomous surveillance, target identification, and situational awareness in contested, EW- and GPS-denied environments. Robust, adversary-resilient models sustain accuracy under attack, ensuring reliable performance in mission-critical operations at the tactical edge.

DetectorGuard: Provably Securing Object Detectors against Localized Patch Hiding Attacks

Research Question: How can object detectors be secured against localized patch hiding attacks using provable methods?

DOI Link GitHub Repository The DetectorGuard paper introduces a framework designed to secure object detectors against localized patch hiding attacks. The approach leverages provable defenses by analyzing the vulnerabilities of object detection models and implementing robust optimization techniques to mitigate adversarial effects. DetectorGuard employs mathematical guarantees to ensure the integrity of detection outputs, even in the presence of adversarial patches. Experimental results demonstrate the effectiveness of the framework in maintaining detection accuracy and robustness across various attack scenarios. The work contributes to advancing the security of computer vision systems in adversarial environments.

For the Army, this research hardens and improves the reliability of object detection for autonomous surveillance, target identification, and situational awareness in contested, EW- and GPS-denied environments. Robust defenses against adversarial attacks preserve operational effectiveness, reduce system vulnerabilities, and sustain trust in mission-critical deployments at the tactical edge.

Learning Inverse Kinodynamics for Accurate High-Speed Off-Road Navigation on Unstructured Terrain

Research Question: How can inverse kinodynamics models improve the accuracy and performance of high-speed off-road navigation on unstructured terrain?

DOI Link This paper introduces a novel approach to high-speed off-road navigation by leveraging inverse kinodynamics models to predict vehicle dynamics and optimize control strategies. The method integrates machine learning techniques to learn terrain-specific dynamics, enabling accurate motion planning and control in unstructured environments. The framework is designed to handle the challenges of high-speed navigation, such as uneven terrain and dynamic obstacles, by providing real-time predictions of vehicle behavior. Experimental results demonstrate significant improvements in navigation accuracy, stability, and speed compared to traditional methods, making the approach suitable for deployment in complex outdoor scenarios.

For the Army, this capability advances high-speed off-road autonomy for reconnaissance, logistics, and rapid response in unstructured, complex terrain. Improved kinodynamic modeling increases maneuver accuracy and stability, elevates operational tempo, and reduces risk to Soldiers in challenging field conditions.

WayFAST: Navigation With Predictive Traversability in the Field

Research Question: How can predictive traversability models improve autonomous navigation in complex field environments?

DOI Link GitHub Repository The WayFast paper introduces a navigation framework that integrates predictive traversability models to enhance autonomous navigation in challenging field environments. The approach combines real-time sensor data with machine learning techniques to predict terrain traversability, enabling robots to make informed decisions about path planning and obstacle avoidance. WayFAST is designed to operate efficiently in dynamic and unstructured environments, leveraging robust algorithms to adapt to changing conditions. Experimental results demonstrate significant improvements in navigation performance, including reduced collision rates and increased traversal efficiency, compared to traditional methods. The framework is particularly suited for deployment in outdoor and off-road scenarios.

For the Army, this capability improves autonomous ground navigation in complex, unstructured terrain for reconnaissance, logistics, and support missions where human access is limited or dangerous. Predictive traversability models increase path safety and efficiency under uncertainty, enhancing mission effectiveness and reducing risk to Soldiers in field operations.

Multi-View Action Recognition using Contrastive Learning

Research Question: How can contrastive learning improve multi-view action recognition in complex environments?

DOI Link GitHub Repository This paper explores the use of contrastive learning for multi-view action recognition, aiming to enhance the ability of models to understand and classify actions from diverse viewpoints. The proposed method leverages contrastive loss to align features across multiple camera perspectives, ensuring robust representation learning. By incorporating multi-view data and optimizing feature consistency, the approach achieves improved recognition accuracy and generalization. Experimental results on benchmark datasets demonstrate significant performance gains compared to traditional methods, highlighting the effectiveness of contrastive learning in handling complex action recognition scenarios.

For the Army, this capability improves analysis and interpretation of human actions across multi-camera surveillance for ISR, security monitoring, and situational awareness in dynamic, contested environments. Enhanced action recognition increases detection accuracy and timeliness at the tactical edge, supporting mission-critical decision-making under uncertainty.

UAV-Assisted Online Machine Learning Over Multi-Tiered Networks: A Hierarchical Nested Personalized Federated Learning Approach

Research Question: How can UAVs enable hierarchical nested personalized federated learning for online machine learning in multi-tiered networks?

DOI Link This paper proposes a hierarchical nested personalized federated learning framework for UAV-assisted online machine learning in multi-tiered networks. The approach leverages UAVs to facilitate data collection and model training across distributed network tiers, ensuring personalized learning for heterogeneous devices. The framework integrates advanced optimization techniques to address challenges such as communication overhead, data privacy, and dynamic network conditions. Experimental results demonstrate improved learning accuracy, reduced latency, and enhanced scalability compared to traditional federated learning methods. The work highlights the potential of UAVs in supporting efficient and adaptive machine learning in complex network environments.

For the Army, this framework employs UAVs to orchestrate distributed, hierarchical machine learning for real-time decision-making, ISR, and battlefield network operations in dynamic, resource-constrained, and communications-contested environments. It delivers robust, adaptive learning with lower latency and resilience to network disruptions, enabling mission-critical analytics and autonomy at the tactical edge.

SSD: A Unified Framework for Self-Supervised Outlier Detection

Research Question: How can self-supervised learning be utilized to create a unified framework for outlier detection?

DOI Link GitHub Repository The SSD paper introduces a unified framework for self-supervised outlier detection that leverages self-supervised learning techniques to identify anomalies in data. The approach uses pretext tasks to learn robust feature representations without requiring labeled data, enabling effective detection of outliers across diverse datasets. SSD integrates feature extraction and anomaly scoring into a single pipeline, optimizing both components for improved performance. Experimental evaluations demonstrate the framework's ability to achieve state-of-the-art results in outlier detection tasks, with high accuracy and scalability across various domains. The work highlights the potential of self-supervised learning in addressing challenges in anomaly detection.

For the Army, this capability strengthens anomaly detection across critical data streams for cybersecurity, ISR, and operational decision-making in contested, low-bandwidth, and edge environments. Robust outlier detection preserves mission system integrity, surfaces emerging threats sooner, and enables proactive mitigation to reduce risk and downtime.

How Does It Feel? Self-Supervised Costmap Learning for Off-Road Vehicle Traversability

Research Question: How can self-supervised learning improve costmap generation for off-road vehicle traversability?

DOI Link GitHub Repository This paper presents a self-supervised learning framework for costmap generation to enhance off-road vehicle traversability. The approach uses sensor data and self-supervised techniques to learn terrain features and predict traversability costs without requiring labeled data. By integrating real-time perception and adaptive learning, the framework generates accurate costmaps that improve path planning and obstacle avoidance in unstructured environments. Experimental results demonstrate significant improvements in navigation performance, including better terrain understanding and reduced traversal errors, compared to traditional methods. The work highlights the potential of self-supervised learning in advancing autonomous navigation in challenging terrains.

For the Army, this capability enhances off-road autonomous navigation for reconnaissance, logistics, and operational support in complex, unstructured, and GPS-denied environments. Improved costmap learning increases path safety and reliability under uncertainty, reducing risk to Soldiers and sustaining mission performance at the tactical edge.

Athena: Probabilistic Verification of Machine Unlearning

Research Question: How can probabilistic methods be used to verify the effectiveness of machine unlearning?

DOI Link GitHub Repository The Athena paper introduces a framework for probabilistic verification of machine unlearning, which ensures that specific data points are effectively removed from machine learning models. The approach uses statistical methods to quantify the likelihood that unlearning has been successfully performed, providing guarantees on data removal without compromising model integrity. Athena integrates probabilistic reasoning and computational techniques to address challenges in verifying unlearning processes across various model architectures. Experimental results demonstrate the framework's ability to provide reliable verification while maintaining efficiency and scalability. The work contributes to advancing privacy-preserving machine learning practices.

For the Army, this capability strengthens data privacy and security in machine learning systems supporting ISR, cybersecurity, and operational decision-making in distributed, sensitive, and contested environments. Probabilistic verification of machine unlearning enforces compliant data removal, reduces retention risk and leakage, and preserves trust in mission-critical models throughout their lifecycle.

Fine-tuning Aligned Language Models Compromises Safety, Even When Users Do Not Intend To!

Research Question: How does fine-tuning aligned language models impact their safety, even in cases where users do not intend to compromise it?

DOI Link GitHub Repository This paper investigates the unintended safety risks associated with fine-tuning aligned language models. The authors demonstrate that even benign fine-tuning processes can lead to misalignment, compromising the model's ability to adhere to safety constraints. The study explores the mechanisms behind these safety failures, including shifts in model behavior and vulnerabilities to adversarial inputs. Experimental results highlight the trade-offs between customization and maintaining alignment, emphasizing the need for robust safeguards during fine-tuning. The findings contribute to understanding the challenges of preserving safety in large-scale language models.

For the Army, this work strengthens the reliability and safety of AI systems used for decision-making, communications, and operational planning in contested, sensitive environments. By enforcing alignment in language models, it prevents unintended behaviors, reduces vulnerabilities and exploitation risk, and preserves mission assurance across edge and enterprise deployments.

PatchGuard: A Provably Robust Defense against Adversarial Patches via Small Receptive Fields and Masking

Research Question: How can small receptive fields and masking be used to create a provably robust defense against adversarial patches?

DOI Link GitHub Repository The PatchGuard paper introduces a defense mechanism against adversarial patch attacks that leverages small receptive fields and masking techniques to ensure robustness. The approach limits the influence of adversarial patches by constraining the receptive field size, preventing malicious alterations from propagating across the model. Additionally, PatchGuard employs masking strategies to isolate and neutralize adversarial regions, ensuring the integrity of predictions. Experimental evaluations demonstrate the effectiveness of PatchGuard in maintaining high accuracy and robustness under various attack scenarios. The framework provides provable guarantees of defense, making it a significant advancement in adversarial robustness research.

For the Army, this capability hardens and improves the reliability of AI systems for autonomous surveillance, target recognition, and decision-making in contested, EW- and GPS-denied environments. Robust defenses against adversarial attacks preserve operational effectiveness, reduce vulnerabilities, and sustain trust in mission-critical deployments at the tactical edge.

VOILA: Visual-Observation-Only Imitation Learning for Autonomous Navigation

Research Question: How can visual-observation-only imitation learning be used to enable autonomous navigation?

DOI Link GitHub Repository The VOILA paper introduces a framework for autonomous navigation that relies solely on visual observation for imitation learning. The approach eliminates the need for explicit action labels by leveraging visual data to train models capable of mimicking expert navigation behaviors. VOILA employs advanced computer vision techniques and neural network architectures to extract meaningful features from visual inputs, enabling robust decision-making in complex environments. Experimental results demonstrate the framework's effectiveness in achieving high navigation accuracy and adaptability across diverse scenarios, highlighting its potential for real-world deployment.

For the Army, this capability enables autonomous navigation from visual observations when traditional sensors are degraded, denied, or impractical, supporting reconnaissance, logistics, and operational support in contested, resource-constrained settings. It reduces sensor burden and SWaP requirements, improving reliability under uncertainty and sustaining mission performance at the tactical edge.

Sigma: Siamese Mamba Network for Multi-Modal Semantic Segmentation

Research Question: How can a Siamese Mamba Network improve multi-modal semantic segmentation for complex environments?

DOI Link GitHub Repository The Sigma paper introduces a Siamese Mamba Network designed for multi-modal semantic segmentation, which integrates data from multiple sensor modalities to enhance segmentation accuracy. The network employs a Siamese architecture to process and align features from different modalities, ensuring robust fusion and improved representation learning. Sigma leverages advanced deep learning techniques to address challenges such as feature misalignment and data heterogeneity. Experimental results demonstrate significant improvements in segmentation performance across diverse datasets, highlighting the framework's adaptability and scalability in complex environments.

For the Army, this capability advances multi-modal data analysis for ISR, autonomous navigation, and battlefield situational awareness in dynamic, contested, and GPS-denied environments. Improved semantic segmentation delivers more accurate sensor fusion and scene understanding, accelerating mission-critical decision-making at the tactical edge.

APPL: Adaptive Planner Parameter Learning

Research Question: How can adaptive learning techniques optimize planner parameters for improved robotic navigation?

DOI Link The APPL paper introduces a framework for adaptive planner parameter learning that optimizes robotic navigation performance in dynamic environments. The approach uses machine learning techniques to adjust planner parameters based on real-time feedback, enabling robots to adapt to changing conditions and improve decision-making. APPL integrates reinforcement learning and optimization methods to enhance path planning efficiency and robustness. Experimental results demonstrate significant improvements in navigation accuracy, computational efficiency, and adaptability across diverse scenarios, highlighting the framework's potential for real-world applications.

For the Army, this capability enhances autonomous navigation for ground robots in dynamic, contested, and GPS-denied environments across reconnaissance, logistics, and operational support. Adaptive planning delivers reliable performance under uncertainty, increasing mission success and resilience at the tactical edge.

Neurotoxin: Durable Backdoors in Federated Learning

Research Question: How can durable backdoors be introduced and maintained in federated learning systems?

DOI Link GitHub Repository The Neurotoxin paper introduces a framework for analyzing and implementing attacks that investigate the vulnerability of federated learning systems to durable backdoor attacks. The approach demonstrates how malicious participants can inject backdoors into the global model during the federated learning process, ensuring persistence across multiple training rounds. The study explores the mechanisms that enable backdoor durability, including model poisoning and adaptive attack strategies. Experimental results highlight the effectiveness of Neurotoxin in maintaining backdoor functionality while evading detection, raising critical concerns about the security of federated learning systems.

For the Army, this work strengthens the security of federated learning across distributed ISR, logistics, and operational planning networks in contested, heterogeneous environments. By identifying and mitigating durable backdoor vulnerabilities, it preserves model integrity and reliability across clients, reduces adversary persistence, and sustains trust in mission-critical AI deployments.

TartanDrive: A Large-Scale Dataset for Learning Off-Road Dynamics Models

Research Question: How can a large-scale dataset like TartanDrive improve the learning of off-road dynamics models for autonomous vehicles?

DOI Link GitHub Repository The TartanDrive paper introduces a large-scale dataset designed to facilitate the learning of off-road dynamics models for autonomous vehicles. The dataset includes diverse terrain types, vehicle dynamics data, and environmental conditions, enabling robust training of machine learning models for off-road navigation. TartanDrive leverages high-resolution sensors and real-world driving scenarios to capture detailed dynamics, providing a comprehensive resource for developing and validating predictive models. Experimental results demonstrate the dataset's utility in improving model accuracy and generalization across challenging off-road environments. The work highlights the importance of large-scale data in advancing autonomous vehicle capabilities.

For the Army, this resource accelerates development of autonomous vehicles for complex, unstructured, and contested off-road environments. High-fidelity dynamics models improve prediction and control, boosting mission success in reconnaissance, logistics, and operational support while reducing risk to Soldiers.

2.7 High-Level Priority Research Questions

During Phase I of A2I2, AI/ML advanced through deep learning, larger datasets, and increased computational capabilities. When the program was formulated in 2018, the emphasis was on foundational techniques (computer vision, autonomy, adversarial robustness). As LLMs matured in 2020–2023, the program incorporated language-grounded planning, open-vocabulary perception, and human–AI interaction.

The funded efforts were guided by a set of high-level research questions:

1. **How can autonomous systems make robust, adaptive decisions in contested, dynamic environments?** (e.g., *WayFAST*, *Learning Inverse Kinodynamics*, *Planning for Manipulation among Movable Objects*)
2. **How can AI models be secured against adversarial attacks and vulnerabilities?** (e.g., *PatchGuard*, *RobustBench*, *DetectorGuard*)
3. **How can self-supervised and unsupervised methods reduce labeled-data dependence while sustaining performance?** (e.g., *SSD: Self-Supervised Outlier Detection*, *Unsupervised Depth Completion*, *Self-Supervised Costmap Learning*)
4. **How can LLMs and vision–language models be integrated into embodied and enterprise pipelines for perception, planning, and decision support?** (e.g., *ProgPrompt*, *LLM-Planner*, *ConceptGraphs*)
5. **How can datasets and benchmarks accelerate model development, evaluation, and reproducibility?** (e.g., *TartanDrive*, *RobustBench*, *Feature 3DGS*)

These questions balance foundational advances with applied objectives, emphasizing operational constraints, data efficiency, robustness, and reproducible evaluation.

3. Key Program Metrics and Army Impact

3.1 Collaboration and Network Impact

The A2I2 served as a hub for collaboration and connectivity between ARL, academia, industry, and federally funded research and development centers (FFRDCs). This section highlights the breadth and depth of A2I2 collaborations across institutions, sectors, and countries through constructing a database of publications and analyzing that data.*

A total of **116** distinct institutions contributed to A2I2 publications. These institutions span academia (78.2% of unique authors), government (11.4%), industry (7.9%), and FFRDCs (2.6%), reflecting the distributed nature of the A2I2 program. Over 90% of unique authors are affiliated with U.S. institutions, while 9.6% are international. Among international collaborators, the largest shares come from Italy (17.9%), Korea (13.4%), and China (11.9%). Other represented countries include Switzerland (10.4%), the United Kingdom (10.4%), Canada (7.5%), India (7.5%), Cameroon (6%), Germany, Israel, the Netherlands, Australia, Austria, Qatar, and Singapore (all less than 5%).†

Table 1 lists the top 15 institutions by number of unique authors, along with the total author appearances across papers and the number of unique papers represented. ARL has the highest number of unique authors and appears on the most papers, underscoring its central role as the program's anchor institution. ARL is followed by leading academic universities, where contributions predominantly represent faculty, postdoctoral researchers, and graduate students.

Figure 1 provides a complementary treemap visualization. Larger blocks represent institutions with greater numbers of unique authors, illustrating the wide distribution of contributors across ARL, major research universities, and industry partners.

Table 2 shifts focus from authors to papers, showing how frequently institutions collaborated with ARL. Some institutions tended to have a closer collaboration with ARL than others, e.g., University of Southern California (USC) had 45 of its 56

*During the preparation of this work, the authors used AskSage to assist in writing code to extract data from APIs, clean data, parse structured data, and prepare it for analysis. Much of the following quantitative analysis stems from creating this database.

†Percentages listed for countries are relative to the subset of international authors; relative to the total set of authors, each percentage is about one-tenth as large.

Table 1. Top 15 institutions by number of unique authors, author appearances, and unique papers.

Institution	Unique Authors	Author Appearances	Unique Papers
DEVCOM Army Research Laboratory	65	185	110
Purdue University	55	180	49
Carnegie Mellon University	50	148	45
University of Southern California	50	156	56
University of California, Los Angeles	45	79	20
The University of Texas at Austin	35	87	25
Texas A&M University	30	48	18
Princeton University	23	116	35
Johns Hopkins University	19	37	13
The Ohio State University	18	31	8
George Mason University	17	57	24
Google	17	17	2
University of Wisconsin–Madison	16	30	9
JHU Applied Physics Laboratory	15	46	12
University of California, Berkeley	14	16	5

papers with ARL coauthors. Interestingly, on average, USC collaborated with one other institution (i.e., 2.1 institutions per paper) and this table shows that often that institution was ARL. Johns Hopkins University (JHU) and JHU Applied Physics Laboratory consistently published with three or more institutions per paper, highlighting strong ties between Johns Hopkins, its affiliated FFRDC, and ARL. On average, ARL papers included 2.7 institutions and 6.2 authors, underscoring the collaborative nature of the program.

Together, these results demonstrate the central role of ARL in the A2I2 program while also emphasizing the program's wide-ranging collaborations with leading universities, industry, and international partners.

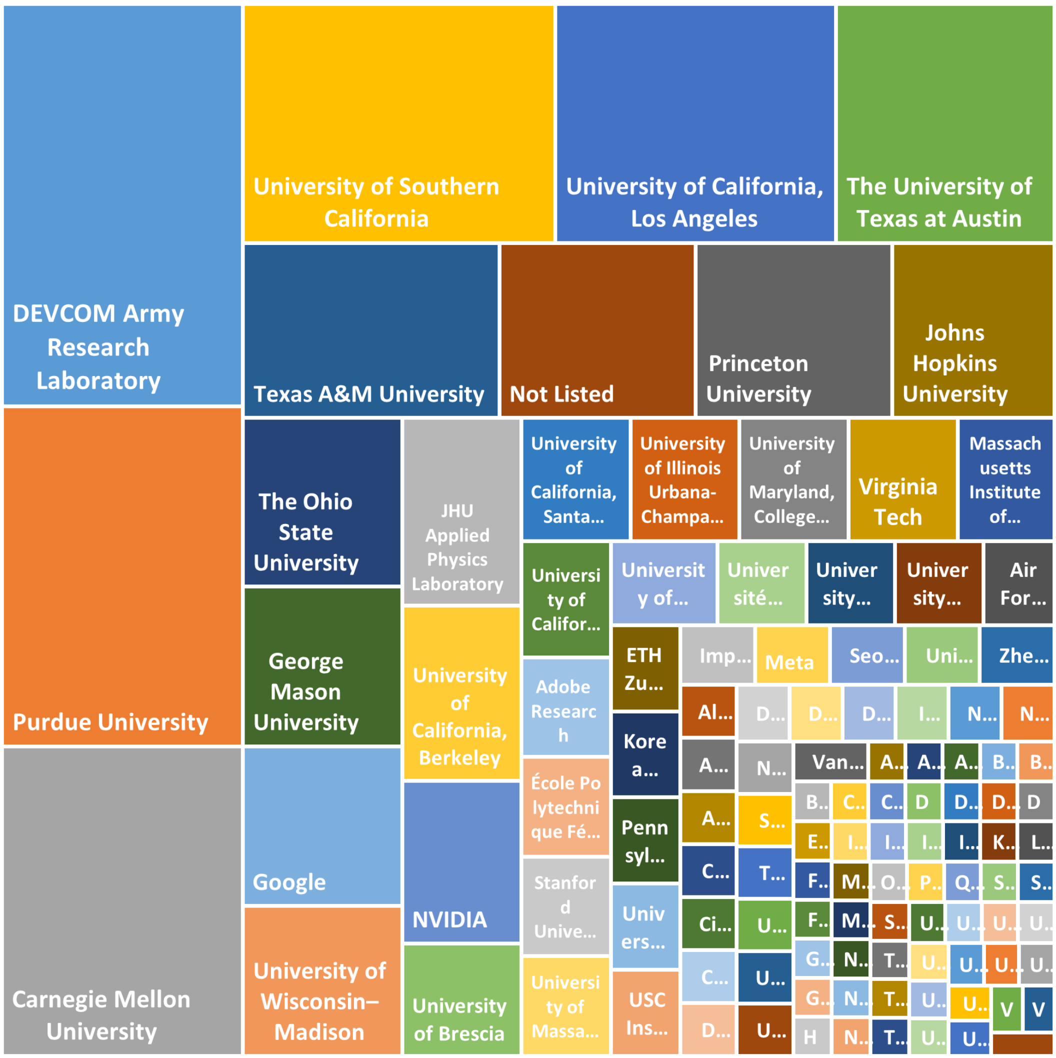

Figure 1. Treemap illustrating the number of unique authors per institution, complementing the data in Table 1.

Table 2. Top non-ARL institutions, with ARL for comparison, ranked by unique papers, ARL collaboration frequency, and average institutions/authors per paper.

Institution	Unique Papers	Papers with ARL	Average Institutions/Paper	Average Authors/Paper
DEVCOM Army Research Laboratory	110	—	2.7	6.2
University of Southern California	56	45	2.1	4.7
The University of Texas at Austin	25	14	2.7	7.1
Johns Hopkins University	13	10	3.3	7.5
JHU Applied Physics Laboratory	12	10	3.5	8.5
Purdue University	49	9	2.1	5.6
University of California, Los Angeles	20	9	2.4	6.8
George Mason University	24	6	2.2	6.6
University of Brescia	7	5	3.3	6.0
Massachusetts Institute of Technology	6	5	4.5	10.8
Carnegie Mellon University	45	4	1.7	4.5
Texas A&M University	18	3	1.4	4.1
University of California, Santa Barbara	6	3	1.8	4.5
University of California, San Diego	7	3	1.6	3.1
University of Dschang	3	3	3.0	5.7
USC Institute for Creative Technologies	3	3	2.7	9.0
The University of Texas at El Paso	3	3	3.0	5.7

3.2 Author and Talent Development

The A2I2 cultivated talent and transferred knowledge between ARL, academia, industry, and FFRDCs. This section highlights the breadth and depth of A2I2 collaborators—all authors and institutions are listed in Appendix B.

The A2I2 portfolio engaged a remarkable **742** unique authors, underscoring the scale of collaboration and the diffusion of knowledge across institutions. Table 3 highlights the most prolific non-ARL authors contributing to the A2I2 portfolio. Several faculty members stand out with sustained and repeated engagement, often serving in senior (author-last) positions that suggest leadership of research groups. For example, Profs. Prateek Mittal (Princeton University) and Viktor Prasanna (USC) contributed to 34 publications, appearing predominantly as the last author (29 and 27, respectively), reflecting their roles as lead investigators.

Table 3. Top non-ARL authors by publications.

Author	**Institution**	**Papers** No. **(1, 2) (n-1, n)**
Prateek Mittal	Princeton University	34 (0, 2) (3, 29)
Viktor K. Prasanna	University of Southern California	34 (0, 1) (5, 27)
Xuesu Xiao	George Mason University	24 (5, 0) (5, 10)
Saurabh Bagchi	Purdue University	22 (0, 4) (8, 8)
Oliver Kroemer	Carnegie Mellon University	17 (0, 7) (2, 15)
Mung Chiang	Purdue University	16 (0, 1) (7, 6)
C.-C. Jay Kuo	University of Southern California	16 (0, 0) (0, 16)
Maxim Likhachev	Carnegie Mellon University	14 (0, 5) (2, 12)
Saeed Mahloujifar	Princeton University	14 (0, 8) (8, 0)
Peter Stone	The University of Texas at Austin	14 (0, 0) (2, 3)
Somali Chaterji	Purdue University	12 (0, 0) (0, 9)
Vikash Sehwag	Princeton University	11 (2, 1) (4, 0)
Christopher G. Brinton	Purdue University	10 (0, 0) (1, 6)
Rama Chellappa	Johns Hopkins University	10 (0, 1) (0, 8)

The distribution of positions also provides insight into career stage and research leadership. Authors such as Prof. C.-C. Jay Kuo (USC) and Prof. Rama Chellappa (JHU) appeared almost exclusively in the last-author role, consistent with their status as senior researchers. In contrast, several authors appear primarily in middle or second-to-last positions, suggesting sustained collaboration within larger teams. Taken together, these results underscore that A2I2 publications were not isolated collaborations but instead drew on leading faculty who contributed repeatedly across the program, often anchoring multi-author, multi-institution efforts and training early-career researchers through coauthorship.

Table 4 highlights the top authors by number of first-author publications. In contrast to Table 3, which emphasizes senior investigators and group leaders, this table reflects the contributions of early-career researchers—graduate students, postdoctoral fellows, junior staff, and assistant professors—who most often appear as first author. These positions typically indicate primary responsibility for the research execution and manuscript preparation, underscoring how the A2I2 program cultivated talent and enabled hands-on experience in high-visibility projects.

Several names stand out with first author papers for performing the research and writing up the findings. For example, Prof. Xuesu Xiao (an A2I2 PI) had the most publications on this list, and these were split between first author publications, second-to-last, and last author publications—a sign of emerging faculty who often have to split time between doing research and advising research. Bingyi Zhang, Hongyu Fu (both University of Southern California), and Arun Reddy (JHU Applied Physics Laboratory) had nine, seven, and seven publications, respectively. While most entries correspond to students or postdocs, the presence of senior researchers such as David A. Handelman (JHU APL) shows that not all first-author roles are restricted to early career researchers. Overall, the distribution of first-author publications demonstrates the depth of talent development within the program: A2I2 enabled junior researchers across multiple institutions to take leading roles in publications, while also engaging senior staff in specialized areas. This balance of mentorship and independent contribution reflects a healthy research ecosystem that both develops the next generation of talent and leverages established expertise.

After highlighting the contributions of students, postdocs, and early-career researchers (Table 4), it is also worth noting the strong presence of ARL scientists in this field. Table 5 lists the most prolific ARL contributors. These include senior researchers such as Rajgopal Kannan, Suya You, and Carl Busart, whose publication records span both early and late positions in author lists, as well as other researchers with targeted contributions. Of note in this table is Aamodh Suresh, who wrote three first-author papers and was advised by Prof. Sònia Martínez

Table 4. Top authors by first author publications.

Author	Institution	Papers No. **(1, 2) (n-1, n)**
Xuesu Xiao	George Mason University	24 (5, 0) (5, 10)
Bingyi Zhang	University of Southern California	9 (5, 2) (0, 0)
Hongyu Fu	University of Southern California	7 (4, 3) (0, 0)
Arun Reddy	JHU Applied Physics Laboratory	7 (4, 0) (0, 0)
Akhil Bandarupalli	Purdue University	6 (4, 2) (0, 0)
Haresh Karnan	The University of Texas at Austin	6 (4, 1) (0, 0)
Xiangyu Qi	Princeton University	6 (4, 1) (0, 0)
Mohit Sharma	Carnegie Mellon University	6 (4, 2) (1, 0)
Chong Xiang	Princeton University	6 (4, 0) (0, 0)
Sihui Dai	Princeton University	5 (4, 0) (0, 0)
Jacky Liang	Carnegie Mellon University	5 (4, 1) (0, 0)
Joshua C. Zhao	Purdue University	5 (4, 0) (0, 0)
Zhiruo Zhou	University of Southern California	5 (4, 1) (0, 0)
Sean Kulinski	Purdue University	4 (4, 0) (0, 0)
Ketul Shah	Johns Hopkins University	7 (3, 2) (0, 0)
Alex LaGrassa	Carnegie Mellon University	6 (3, 1) (1, 0)
Ashwinee Panda	Princeton University	6 (3, 2) (0, 0)
Matthew Sivaprakasam	Carnegie Mellon University	6 (3, 1) (0, 0)
Aniket Datar	George Mason University	5 (3, 0) (1, 0)
David A. Handelman	JHU Applied Physics Laboratory	5 (3, 0) (1, 0)
Shijie Zhou	University of California, Los Angeles	5 (3, 1) (0, 0)
Mustafa Abdallah	Purdue University	4 (3, 1) (0, 0)
Shohin Mukherjee	Carnegie Mellon University	4 (3, 1) (1, 0)
Dhruv Parikh	University of Southern California	4 (3, 1) (0, 0)
Sasindu Wijeratne	University of Southern California	4 (3, 1) (0, 0)
Ce Zhang	Carnegie Mellon University	4 (3, 0) (0, 0)
Sophia J. Abraham	University of Notre Dame	3 (3, 0) (0, 0)
Hong-Shuo Chen	University of Southern California	3 (3, 0) (0, 0)
Tianqi Li	Texas A&M University	3 (3, 0) (0, 0)
Dhruv Mauria Saxena	Carnegie Mellon University	3 (3, 0) (0, 0)
Xinyu Tang	Princeton University	3 (3, 0) (0, 0)

at UCSD, and became an ARL postdoctoral researcher, showing how the A2I2 program also helped as a recruiting tool for ARL.

Table 5. Top ARL authors by publications.

Author	**Institution**	**Papers** No. **(1, 2) (n-1, n)**
Rajgopal Kannan	DEVCOM Army Research Laboratory	27 (0, 12) (14, 0)
Suya You	DEVCOM Army Research Laboratory	19 (0, 2) (10, 1)
Carl Busart	DEVCOM Army Research Laboratory	11 (0, 0) (6, 5)
Celso Miguel de Melo	DEVCOM Army Research Laboratory	10 (0, 1) (5, 1)
Garrett Warnell	DEVCOM Army Research Laboratory	9 (0, 2) (1, 0)
Lance Kaplan	DEVCOM Army Research Laboratory	8 (0, 1) (2, 2)
Vinod K. Mishra	DEVCOM Army Research Laboratory	5 (0, 1) (5, 0)
Shuowen Hu	DEVCOM Army Research Laboratory	4 (0, 1) (0, 0)
John G. Rogers	DEVCOM Army Research Laboratory	4 (0, 0) (1, 0)
Aamodh Suresh	DEVCOM Army Research Laboratory	3 (3, 0) (0, 0)
Damon M. Conover	DEVCOM Army Research Laboratory	3 (0, 2) (1, 0)
Venkateswara Dasari	DEVCOM Army Research Laboratory	3 (0, 1) (0, 0)
Charles Kamhoua	DEVCOM Army Research Laboratory	3 (0, 0) (0, 3)
Heesung Kwon	DEVCOM Army Research Laboratory	3 (0, 0) (2, 0)
Kimberly A. Pollard	DEVCOM Army Research Laboratory	3 (0, 2) (1, 0)
Noah Weston	DEVCOM Army Research Laboratory	3 (0, 0) (1, 0)

Table 5 highlights the most prolific authors within ARL for the A2I2 program. To put these contributions in context, Figure 2 shows the proportion of papers with ARL authorship. Roughly two-thirds of the papers (64%) had no ARL authors, while 36% included at least one ARL contributor. Of these, the majority had a single ARL author (22%), with a smaller fraction involving two (14%) or more ARL researchers. This pattern underscores that ARL participation is both frequent and collaborative, with most contributions embedded in larger multi-institution teams rather than concentrated within ARL-only groups.

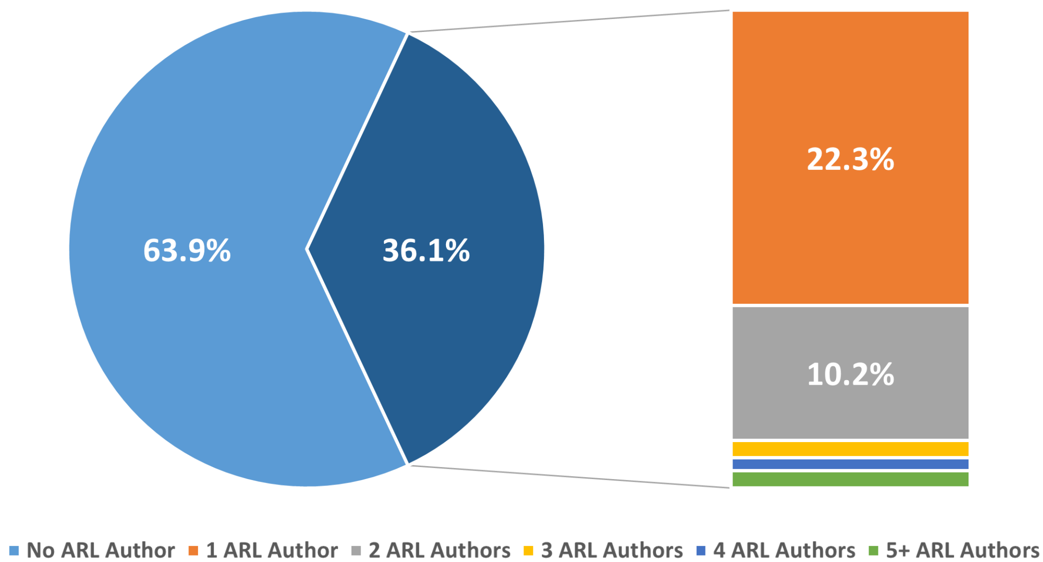

Figure 2. Proportion of publications with ARL authors. About one-third of papers include at least one ARL author, most often as a single contributor alongside external collaborators.

3.3 Dissemination and Visibility

The A2I2 portfolio demonstrated strong dissemination and visibility across the research community. In total, **311 publications** were identified, with the overwhelming majority (**98.1%**) carrying a DOI, ensuring they are trackable and citable (Table 6). Dissemination occurred across a wide range of venues, with IEEE Xplore as the most common outlet (59.8%), followed by arXiv (16.7%), the ACM Digital Library (6.8%), and other prominent publishers such as SPIE, Elsevier, AAAI, and Springer Nature. Notably, arXiv play a distinctive role: many are precursors to conference or journal papers, while others serve as standalone contributions that accelerate dissemination, accrue citations, and broaden access. Several A2I2-related arXiv papers achieved exceptional impact, including works in the top 1% and 10% most cited in their fields, with one exceeding 100 citations.

Table 6. Publication venues for dissemination.

Category	Papers	Percentage
Has DOI	305	98.1%
Does not have DOI	6	1.9%
IEEE Xplore	186	59.8%
arXiv	52	16.7%
ACM digital library	21	6.8%
Elsevier	9	2.9%
SPIE	8	2.6%
AAAI	7	2.3%
Springer Nature	4	1.3%
IEEE International Conference on Robotics and Automation (ICRA)	37	11.9%
IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)	21	6.8%
IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)	11	3.5%
IEEE High Performance Extreme Computing Conference (HPEC)	8	2.6%
IEEE/CVF International Conference on Computer Vision (ICCV)	5	1.6%
AAAI International Symposium on Combinatorial Search	4	1.3%
IEEE Conference on Decision and Control (CDC)	4	1.3%
IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)	4	1.3%
SPIE Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications	3	1.0%
IEEE International Conference on Automatic Face and Gesture Recognition	3	1.0%
IEEE International Conference on Image Processing (ICIP)	3	1.0%
IEEE International Parallel and Distributed Processing Symposium (IPDPS)	3	1.0%
IEEE Symposium on Security and Privacy	3	1.0%
IEEE International Workshop on Technologies for Defense and Security (TechDefense)	3	1.0%

High-profile conferences featured prominently in the dissemination strategy. These included the IEEE International Conference on Robotics and Automation (ICRA, 11.9%), the Conference on Computer Vision and Pattern Recognition (CVPR, 6.8%), the International Conference on Intelligent Robots and Systems (IROS, 3.5%), and the High Performance Extreme Computing Conference (HPEC, 2.6%), among others. Papers also appeared in leading venues in security and privacy, robotics and autonomy, computer vision, and defense-focused forums (e.g., SPIE AI/ML for MDO, TechDefense). This mix reflects a deliberate balance between top-tier technical venues that build global credibility and domain-specific conferences that directly engage Army and DoD stakeholders. Presenting at these venues also develops the next generation of researchers by exposing them to peer scrutiny, networking opportunities, and professional visibility. Beyond publications, A2I2 teams disseminated their work through poster competitions, workshops, VIP site visits, and other outreach activities that amplified impact.

Acknowledgment of ARL and A2I2 support was widespread but not universal (Figure 3). Approximately 69.5% of publications cited A2I2 funding with the correct agreement number, while 21.2% did not include any acknowledgments, and 9.3% mentioned A2I2 without citing the agreement number. While the majority of authors credited Army sponsorship, there remains room to improve consistency in acknowledgment practices to maximize visibility of Army investment.

By publishing in world-class venues, ensuring broad DOI coverage, and engaging across both academic and defense-focused platforms, A2I2 researchers positioned Army-sponsored work at the forefront of global conversations in artificial intelligence, robotics, autonomy, and security. This strategy advanced scientific credibility, operational relevance, and talent development, reinforcing A2I2's central mission: to cultivate talent, transfer knowledge, and strengthen the Army's scientific and technological edge.

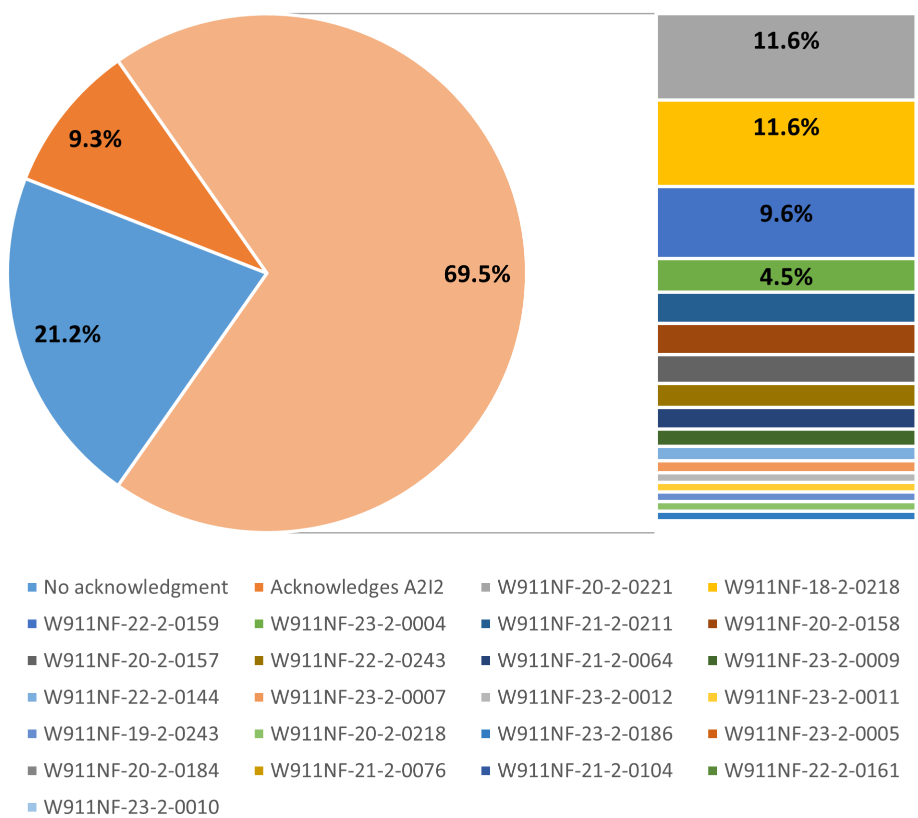

Figure 3. Proportion of publications that acknowledge A2I2 funding and cite the agreement number (breakout), that only acknowledge A2I2 funding, and that do not acknowledge funding from ARL.

3.4 Quality and Reach

Citations remain one of the most widely recognized indicators of scientific impact and quality. The total number of citations of the A2I2 works is **4266** citations at the time of this report (as of 3 DEC 2025). While not a perfect measure, they serve as an external signal of how much a work has been integrated into and influenced subsequent research. In the context of open science, where dissemination and accessibility are key, citation counts help demonstrate both the uptake of research outputs and their contribution to broader scientific progress.

However, citations are inherently a lagging metric. It can take months to years for a publication to be read, applied, and cited by others. As such, the current citation record often underestimates the eventual influence of more recent works. This temporal delay must be considered when evaluating the immediate impact of open science outputs, particularly for newer publications.

To help visualize how scholarly influence has accumulated over time, Figure 4 presents a waterfall plot tracing the year-by-year growth in citations to A2I2-supported publications. First, each publication within the top 30 is represented by a different color to show their citation evolution; the publications are ordered in terms of total number of citations, with the highest number on the bottom of the plot. Note that any publication not in the top 30 was grouped into an "Other" category. While the top 30 accounts for a large number of citations, the other publications also are increasingly being cited by the community. Overall, the total citation activity begins to rise modestly in 2022, accelerates sharply in 2023, and continues to expand through 2024–2025 as more works gain traction within the scientific community. Although citation dynamics typically follow a long-tail lifespan, the higher citation totals observed in 2025—exceeding those from 2024—suggest that recently published A2I2 research is actively driving new scientific progress. This pattern highlights both the continued visibility of earlier publications and the strong uptake of newly emerging results.

Table 7 lists the top 30 most-cited publications in the dataset. The choice of 30 corresponds to the h-index of this corpus, i.e., 30 publications have been cited at least 30 times. These top-cited works collectively account for 2,633 citations, representing 61.7% of the total citation volume. Individually, these publications range from 32 to 446 citations—nearly an order of magnitude—which underscores both

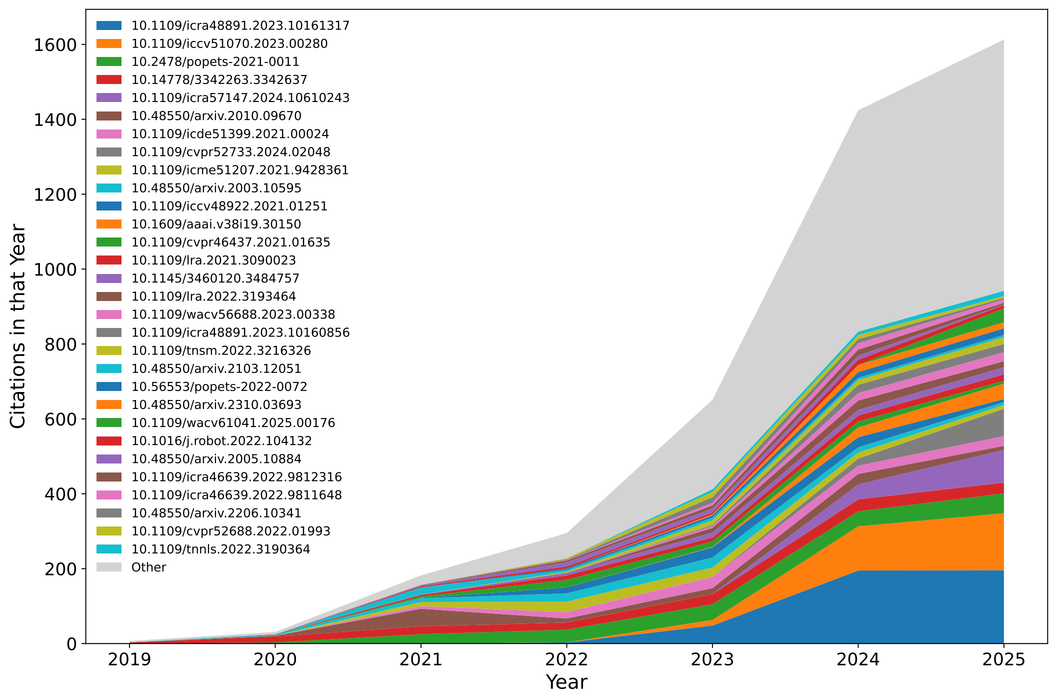

Figure 4. Year-to-year accumulation of citations to A2I2-supported publications, illustrating the accelerating growth in scholarly uptake from 2022 through 2025.

the heterogeneity of scientific uptake and the disproportionate influence of a small set of high-impact papers. The complete citation dataset is provided in Appendix C. Together, these results demonstrate that A2I2-enabled research has achieved substantial reach, with key contributions gaining broad and sustained visibility even as others continue to accumulate citations more gradually.

Table 7. Most cited A2I2 publications (as of 3 DEC 2025).

1. ProgPrompt: Generating Situated Robot Task Plans using Large Language Models (2023)—446 citations
2. LLM-Planner: Few-Shot Grounded Planning for Embodied Agents with Large Language Models (2023)—293 citations
3. Falcon: Honest-Majority Maliciously Secure Framework for Private Deep Learning (2021)—199 citations
4. Efficient task-specific data valuation for nearest neighbor algorithms (2019)—148 citations
5. ConceptGraphs: Open-Vocabulary 3D Scene Graphs for Perception and Planning (2024)—132 citations
6. RobustBench: a standardized adversarial robustness benchmark (2020)—116 citations
7. Enabling Efficient Cyber Threat Hunting With Cyber Threat Intelligence (2021)—102 citations
8. Feature 3DGS: Supercharging 3D Gaussian Splatting to Enable Distilled Feature Fields (2024)—94 citations
9. DefakeHop: A Light-Weight High-Performance Deepfake Detector (2021)—90 citations
10. Systematic Evaluation of Privacy Risks of Machine Learning Models (2020)—85 citations
11. Unsupervised Depth Completion with Calibrated Backprojection Layers (2021)—79 citations
12. Visual Adversarial Examples Jailbreak Aligned Large Language Models (2024)—67 citations
13. Robust and Accurate Object Detection via Adversarial Learning (2021)—65 citations

14. Learning Inverse Kinodynamics for Accurate High-Speed Off-Road Navigation on Unstructured Terrain (2021)—60 citations
15. DetectorGuard: Provably Securing Object Detectors against Localized Patch Hiding Attacks (2021)—55 citations
16. WayFAST: Navigation With Predictive Traversability in the Field (2022)—55 citations
17. Multi-View Action Recognition using Contrastive Learning (2023)—48 citations
18. How Does It Feel? Self-Supervised Costmap Learning for Off-Road Vehicle Traversability (2023)—47 citations
19. UAV-Assisted Online Machine Learning Over Multi-Tiered Networks: A Hierarchical Nested Personalized Federated Learning Approach (2022)—46 citations
20. Athena: Probabilistic Verification of Machine Unlearning (2022)—44 citations
21. SSD: A Unified Framework for Self-Supervised Outlier Detection (2021)—44 citations
22. Sigma: Siamese Mamba Network for Multi-Modal Semantic Segmentation (2025)—40 citations
23. Fine-tuning Aligned Language Models Compromises Safety, Even When Users Do Not Intend To! (2023)—38 citations
24. PatchGuard: A Provably Robust Defense against Adversarial Patches via Small Receptive Fields and Masking (2020)—36 citations
25. APPL: Adaptive Planner Parameter Learning (2022)—36 citations
26. VOILA: Visual-Observation-Only Imitation Learning for Autonomous Navigation (2022)—35 citations
27. Synthetic Generation of Face Videos with Plethysmograph Physiology (2022)—34 citations

28. TartanDrive: A Large-Scale Dataset for Learning Off-Road Dynamics Models (2022)—34 citations

29. Neurotoxin: Durable Backdoors in Federated Learning (2022)—33 citations

30. Incorporating physics into data-driven computer vision (2023)—32 citations

Finally, to visualize the structure of how these citations flow into and within the program, Figure 5 maps each original publication with at least one citation to sequential numeric identifiers and links them to the unique identifiers of all works that cite them. Each point in the scatter plot represents a citing–cited pair, with original publication IDs on the x-axis and citing-work IDs on the y-axis. Of the 4,266 total citations shown, 4,109 originate from external publications, demonstrating the breadth of scientific impact beyond A2I2. Meanwhile, 157 internal citations, spread across 80 A2I2 papers, indicate that some research lines within the program build directly on earlier A2I2 findings. Vertically clustered points identify individual citing publications that reference multiple A2I2 works, whereas horizontal clusters reveal A2I2 publications repeatedly cited across different external outputs. Overall, this citation-mapping visualization highlights the dense and growing network of scholarly influence generated by A2I2-supported research.

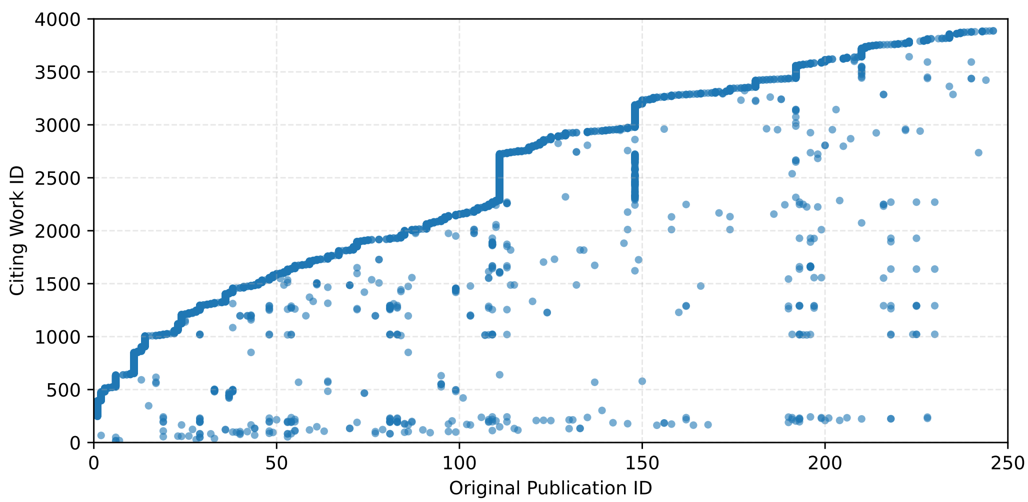

Figure 5. Citation relationships between the A2I2-related publications and all works that reference them.

To complement citation-based measures of quality and reach, the aggregated concept scores provide a quantitative "fingerprint" of the scientific domains represented within the A2I2 publication portfolio. Table 8 lists the top 25 concepts, ranked by the summed relevance scores across all publications. As expected, broad level-0 and level-1 domains—such as computer science, artificial intelligence, machine learning, and computer vision—dominate the ranking. This pattern arises from the structure of the OpenAlex taxonomy: when a publication is tagged with a more specific level-2 or level-3 concept, its higher-level ancestors are also included (e.g., a paper associated with robotics or terrain reasoning contributes to the overarching computer science category). As a result, general domains naturally accumulate larger total scores, while more specific concepts such as terrain, motion planning, object detection, and robustness appear with lower absolute values yet signal the focused areas of technical depth within the portfolio. The choice of relevance-score threshold influences the distribution—stricter cutoffs (e.g., 0.5 or 0.7) prune many level-2 and level-3 concepts—but the overall pattern remains stable: research supported by A2I2 spans a broad computer-science foundation with concentrated strengths in AI, ML, robotics, perception, terrain reasoning, and robustness.

Table 8. Top 25 aggregated concepts across A2I2 publications.

Rank	Concept	Level	Aggregate Score
1	Computer science	0	212.68
2	Artificial intelligence	1	115.79
3	Machine learning	1	39.35
4	Computer vision	1	31.76
5	Robot	2	23.92
6	Task (project management)	2	17.49
7	Mathematics	0	16.43
8	Robustness (evolution)	3	15.86
9	Pattern recognition (psychology)	2	15.49
10	Human–computer interaction	1	15.12
11	Artificial neural network	2	14.69
12	Computer security	1	13.97
13	Adversarial system	2	12.86
14	Algorithm	1	12.48
15	Object (grammar)	2	12.42
16	Terrain	2	11.94
17	Distributed computing	1	11.47
18	Inference	2	10.92
19	Reinforcement learning	2	10.64
20	Theoretical computer science	1	9.78
21	Data mining	1	9.76
22	Motion planning	3	8.71
23	Benchmark (surveying)	2	8.52
24	Object detection	3	8.37
25	Graph	2	8.06

4. Conclusion and Recommendations

The future of the A2I2 lies in coupling large pretrained models (LPTMs) and multimodal models with Army capability gaps.These models offer transformative potential for all military capabilities: electronic warfare, electromagnetic spectrum operations, C2 and counter-command and control, accelerated decision-making, and autonomous systems. To ensure the program remains aligned with Army priorities while preserving the exploratory nature of basic research, the following recommendations are proposed:

1. **Focus Basic Research on Foundational Challenges in LPTMs and Multimodal Models.** Basic research should address core scientific questions that underpin LPTMs and their applications to Army-relevant challenges. Key areas of focus include:

   - Developing custom multimodal models that integrate text, vision, and other modalities for open-world reasoning and action.
   - Advancing reasoning capabilities, such as in-context learning, chain-of-thought reasoning, and long-term planning, to support autonomous agents and decision aids.
   - Exploring digital twin and open-world models for robotics and autonomous systems in Army-relevant environments (e.g., logistics, reconnaissance, search-and-rescue).
   - Investigating methods for knowledge distillation and model compression to deploy smaller, efficient models at the edge.
   - Enabling adaptation and continual learning to ensure LPTMs remain relevant as operational needs and data evolve.

2. **Leverage ARL Scientists as the Bridge Between Basic Research and Army Needs.** ARL scientists should serve as the primary conduit for translating basic research outcomes into applied research opportunities and into the Army ecosystem. This approach ensures that academic researchers remain focused on fundamental discoveries while ARL scientists identify and communicate potential applications to Army stakeholders. ARL's role as a collaborator in this research is essential for aligning research with modernization priorities.

3. **Explore Capstone Demonstrations to Bridge Basic and Applied Research.** Where appropriate, basic research efforts should culminate in capstone demonstrations or science experiments that explore potential applications of LPTMs. These demonstrations can provide insights into how foundational discoveries translate into operational capabilities, inform future pilot efforts, and identify gaps for further research. Such efforts should remain exploratory and avoid turning basic research into near-term engineering projects.

4. **Balance Open Science and Controlled Dissemination to Maximize Transition Potential.** To accelerate the transition of research outputs while safeguarding sensitive information, A2I2 should adopt a balanced approach to dissemination:

    - Publish research findings on open platforms like arXiv to ensure broad accessibility and reproducibility for non-sensitive research outputs. Open dissemination fosters collaboration and accelerates scientific progress across the broader research community.
    - For research outputs that involve sensitive or mission-critical information, collaboration with ARL and other Army partners is essential to ensure secure handling and ingestion into Army environments. Controlled dissemination mechanisms should be established to protect sensitive data while enabling its use within the Army.
    - Utilize ARL's research protection and dissemination strategy for each research output based on its sensitivity, operational relevance, and potential impact. This ensures that open science practices are applied judiciously while maintaining security and alignment with Army priorities.
    - Use secure platforms to facilitate collaboration and controlled sharing of sensitive research outputs among Army researchers, academic partners, and industry stakeholders.

5. **Build a Research Ecosystem That Leverages Academia, Industry, and DOW Strengths.** Given the high barriers to entry for LPTM research (e.g., data, compute, talent), A2I2 should foster a collaborative ecosystem:

    - Strengthen partnerships with academia and industry through collaborative agreements focused on fundamental research topics.

- Establish agreements with industry to gain early access to cutting-edge models and tools.
- Collaborate with other service labs and DOW agencies to align research efforts and share resources.

6. **Incorporate Responsible AI Principles into LPTM Research.** In line with DOW guidance on responsible AI, A2I2 should integrate risk mitigation into its research programs. This includes:
   - Developing benchmarks and maturity models to evaluate LPTMs for Army use cases.
   - Addressing challenges such as hallucinations, data provenance, and model misalignment.
   - Ensuring that LPTMs are deployed securely and ethically, with appropriate guardrails to prevent misuse.

By focusing on these recommendations, A2I2 can position itself as a leader in AI research, advancing the state of the art while addressing Army capability gaps. This approach balances the exploratory nature of basic research with the need to deliver impactful outcomes for Army modernization.

Appendix A. Army Artificial Intelligence Innovation Institute (A2I2) Research Publications

A total of 311 publications were identified as stemming from A2I2-supported research, of which 305 included Digital Object Identifiers (DOIs). Publications were collected and verified through multiple methods. First, the Dimensions platform within DTIC was used to search for the cooperative agreement number in both standard (e.g., "W911NF-18-2-0218") and unhyphenated (e.g., "W911NF1820218") formats, yielding publications that explicitly cited the agreement in the acknowledgments section. Second, searches using the term "A2I2" identified additional publications where authors acknowledged A2I2 funding without listing the agreement number. Finally, principal investigators provided comprehensive lists of all publications produced under A2I2 funding, which were then cross-referenced with findings from the above searches. This ensured capture of papers that were not otherwise retrievable due to variations in publication formats (e.g., proceedings without acknowledgments) or omission of the funding source by the authors.

Of the 311 publications, 216 (70%) explicitly acknowledged U.S. Army Combat Capabilities Development Command (DEVCOM) Army Research Laboratory (ARL) A2I2 support and listed the cooperative agreement number, while an additional 29 (9%) acknowledged A2I2 without listing the agreement number. Together, this accounts for 79% of the identified publications. The remaining 66 publications (21%) did not reference A2I2 or the cooperative agreement number, though some cited other ARL or Army Research Office (ARO) funding sources. Any publications listed as "in preparation" or "in submission" but not verifiable were removed, as were listed excerpts from reports, press releases, or videos, which, while noteworthy, were not peer-reviewed papers.

While most publications are peer-reviewed journal articles or conference proceedings, 52 papers (17%) were published to arXiv, the largest single venue for these works. These were retained because, in many cases, authors posted their conference papers to arXiv for dissemination when proceedings were not published online. Even for those not associated with a conference, some of these arXiv papers have garnered tens to hundreds of citations, demonstrating their quality and impact within the open science research community.

Here are the 305 papers with DOI numbers, in order of publication date.

1. Jia R et al. Efficient task-specific data valuation for nearest neighbor algorithms. Proceedings of the VLDB Endowment 2019; 12(11):1610–1623.

2. Song L, Mittal P. Systematic Evaluation of Privacy Risks of Machine Learning Models. arXiv 2020; 2003.10595.

3. Xiang C, Bhagoji A N, Sehwag V, Mittal P. PatchGuard: A Provably Robust Defense against Adversarial Patches via Small Receptive Fields and Masking. arXiv 2020; 2005.10884.

4. Liang J, Saxena S, Kroemer O. Learning Active Task-Oriented Exploration Policies for Bridging the Sim-to-Real Gap. arXiv 2020; 2006.01952.

5. Azam S S et al. Can we Generalize and Distribute Private Representation Learning?. arXiv 2020; 2010.01792.

6. Croce F et al. RobustBench: a standardized adversarial robustness benchmark. arXiv 2020; 2010.09670.

7. LaGrassa A, Lee S, Kroemer O. Learning Skills to Patch Plans Based on Inaccurate Models. 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) 2020; 9441–9448.

8. Agarwal A, Han Y, Likhachev M. PERCH 2.0 : Fast and Accurate GPU-based Perception via Search for Object Pose Estimation. 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) 2020.

9. Sharma M, Liang J, Zhao J, LaGrassa A, Kroemer O. Learning to Compose Hierarchical Object-Centric Controllers for Robotic Manipulation. arXiv 2020; 2011.04627.

10. Sharma M, Kroemer O. Relational Learning for Skill Preconditions. arXiv 2020; 2012.01693.

11. Lin R, You S, Rao R, Kuo C J. On Relationship of Multilayer Perceptrons and Piecewise Polynomial Approximators. IEEE Signal Processing Letters 2021; 28:1813–1817.

12. Mehta D, Chen T, Tang T, Hauenstein J D. The Loss Surface Of Deep Linear Networks Viewed Through The Algebraic Geometry Lens. IEEE Transactions on Pattern Analysis and Machine Intelligence 2021; 44(9):1.

13. Wagh S et al. Falcon: Honest-Majority Maliciously Secure Framework for Private Deep Learning. Proceedings on Privacy Enhancing Technologies 2021.

14. Sehwag V, Chiang M, Mittal P. SSD: A Unified Framework for Self-Supervised Outlier Detection. arXiv 2021; 2103.12051.

15. Gao P et al. Enabling Efficient Cyber Threat Hunting With Cyber Threat Intelligence. 2021 IEEE 37th International Conference on Data Engineering (ICDE) 2021.

16. Wang R, Wang X, Inouye D I. Shapley Explanation Networks. arXiv 2021; 2104.02297.

17. Zhou Z, Gong Z, Ravikumar P, Inouye D I. Iterative Alignment Flows. arXiv 2021; 2104.07232.

18. Bhagoji A N, Cullina D, Sehwag V, Mittal P. Lower Bounds on Cross-Entropy Loss in the Presence of Test-time Adversaries. arXiv 2021; 2104.08382.

19. Sehwag V et al. Robust Learning Meets Generative Models: Can Proxy Distributions Improve Adversarial Robustness?. arXiv 2021; 2104.09425.

20. Abdallah M, Mitra S, Sundaram S, Bagchi S. HIOA-CPS: Combining Hybrid Input-Output Automaton and Game Theory for Security Modeling of Cyber-Physical Systems. 2021 IEEE Security and Privacy Workshops (SPW) 2021; 49:253–259.

21. Vemula A, Bagnell J A, Likhachev M. CMAX++ : Leveraging Experience in Planning and Execution using Inaccurate Models. Proceedings of the AAAI Conference on Artificial Intelligence 2021; 35(7):6147–6155.

22. Sivaprakasam M, Triest S, Wang W, Yin P, Scherer S. Improving Off-road Planning Techniques with Learned Costs from Physical Interactions. 2021 IEEE International Conference on Robotics and Automation (ICRA) 2021.

23. Liang J, Kroemer O. Contact Localization for Robot Arms in Motion without Torque Sensing. 2021 IEEE International Conference on Robotics and Automation (ICRA) 2021; 2008:6322–6328.

24. Sharma M, Kroemer O. Generalizing Object-Centric Task-Axes Controllers using Keypoints. 2021 IEEE International Conference on Robotics and Automation (ICRA) 2021; 7548–7554.

25. Lee T E, Zhao J, Sawhney A S, Girdhar S, Kroemer O. Causal Reasoning in Simulation for Structure and Transfer Learning of Robot Manipulation Policies. 2021 IEEE International Conference on Robotics and Automation (ICRA) 2021; 4776–4782.

26. Bei X, Yang Y, Soatto S. Learning Semantic-Aware Dynamics for Video Prediction. 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) 2021; 902–912.

27. Chen X et al. Robust and Accurate Object Detection via Adversarial Learning. 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) 2021; 16617–16626.

28. Chen H et al. DefakeHop: A Light-Weight High-Performance Deepfake Detector. 2021 IEEE International Conference on Multimedia and Expo (ICME) 2021; 1–6.

29. Rouhsedaghat M, Wang Y, Hu S, You S, Kuo C J. Low-resolution face recognition in resource-constrained environments. Pattern Recognition Letters 2021; 149:193–199.

30. Lee J et al. Benchmarking Video Object Detection Systems on Embedded Devices under Resource Contention. Proceedings of the 5th International Workshop on Embedded and Mobile Deep Learning 2021; 19–24.

31. Bandarupalli A, Swarup D, Weston N, Chaterji S. Persistent Airborne Surveillance using Semi-Autonomous Drone swarms. Proceedings of the 7th Workshop on Micro Aerial Vehicle Networks, Systems, and Applications 2021; 19–24.

32. Islam F, Salzman O, Agarwal A, Likhachev M. Provably constant-time planning and replanning for real-time grasping objects off a conveyor belt. The International Journal of Robotics Research 2021; 40(12-14):1370–1384.

33. Zou Y, Kim K T, Lin X, Chiang M. Minimizing Age-of-Information in Heterogeneous Multi-Channel Systems. Proceedings of the Twenty-second International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing 2021; 11–20.

34. Xiao X, Biswas J, Stone P. Learning Inverse Kinodynamics for Accurate High-Speed Off-Road Navigation on Unstructured Terrain. IEEE Robotics and Automation Letters 2021; 6(3):6054–6060.

35. Reddyvari V R, Bobbili S C, Parag P, Shakkottai S. Mode-Suppression: A Simple, Stable and Scalable Chunk-Sharing Algorithm for P2P Networks. IEEE/ACM Transactions on Networking 2021; 29(6):2548–2559.

36. Kulinski S, Bagchi S, Inouye D I. Feature Shift Detection: Localizing Which Features Have Shifted via Conditional Distribution Tests. arXiv 2021; 2107.06929.

37. Li T, Krakow L W, Gopalswamy S. Optimizing Consensus-based Multi-target Tracking with Multiagent Rollout Control Policies. 2021 IEEE Conference on Control Technology and Applications (CCTA) 2021; 131–137.

38. Xiang C, Mahloujifar S, Mittal P. PatchCleanser: Certifiably Robust Defense against Adversarial Patches for Any Image Classifier. arXiv 2021; 2108.09135.

39. Wood S et al. Developing Hierarchies for Image Classification Model Evaluation. 2021 4th International Conference on Artificial Intelligence for Industries (AI4I) 2021; 7:34–37.

40. Wong A, Soatto S. Unsupervised Depth Completion with Calibrated Backprojection Layers. 2021 IEEE/CVF International Conference on Computer Vision (ICCV) 2021; 12727–12736.

41. Wang R, Chen X, Cheng M, Tang X, Hsieh C. RANK-NOSH: Efficient Predictor-Based Architecture Search via Non-Uniform Successive Halving. 2021 IEEE/CVF International Conference on Computer Vision (ICCV) 2021; 33:10357–10366.

42. Tang X et al. Mitigating Membership Inference Attacks by Self-Distillation Through a Novel Ensemble Architecture. arXiv 2021; 2110.08324.

43. Xiang C, Mittal P. DetectorGuard: Provably Securing Object Detectors against Localized Patch Hiding Attacks. Proceedings of the 2021 ACM SIGSAC Conference on Computer and Communications Security 2021; 3177–3196.

44. Nguyen H T, Morabito R, Kim K T, Chiang M. On-the-fly Resource-Aware Model Aggregation for Federated Learning in Heterogeneous Edge. 2021 IEEE Global Communications Conference (GLOBECOM) 2021; 14:1–6.

45. Zhou Z, Fu H, You S, Borel-Donohue C C, Kuo C J. UHP-SOT: An Unsupervised High-Performance Single Object Tracker. 2021 International Conference on Visual Communications and Image Processing (VCIP) 2021; 126:1–5.

46. Panda A, Mahloujifar S, Bhagoji A N, Chakraborty S, Mittal P. SparseFed: Mitigating Model Poisoning Attacks in Federated Learning with Sparsification. arXiv 2021; 2112.06274.

47. Zhou Z, Fu H, You S, Kuo C J. UHP-SOT++: An Unsupervised Lightweight Single Object Tracker. APSIPA Transactions on Signal and Information Processing 2022; 11(1).

48. Chen H, Hu S, You S, Kuo C J. DefakeHop++: An Enhanced Lightweight Deepfake Detector. APSIPA Transactions on Signal and Information Processing 2022; 11(2).

49. Mukherjee S, Aine S, Likhachev M. MPLP: Massively Parallelized Lazy Planning. IEEE Robotics and Automation Letters 2022; 7(3):6067–6074.

50. Xu R et al. LiteReconfig. Proceedings of the Seventeenth European Conference on Computer Systems 2022; 334–351.

51. Cody T, Lanus E, Doyle D, Freeman L. Systematic Training and Testing for Machine Learning Using Combinatorial Interaction Testing. 2022 IEEE International Conference on Software Testing, Verification and Validation Workshops (ICSTW) 2022; 102–109.

52. Dai S, Mahloujifar S, Mittal P. Formulating Robustness Against Unforeseen Attacks. arXiv 2022; 2204.13779.

53. Xiao X et al. APPL: Adaptive Planner Parameter Learning. Robotics and Autonomous Systems 2022; 154:104132.

54. Dai S, Mahloujifar S, Mittal P. Parameterizing Activation Functions for Adversarial Robustness. 2022 IEEE Security and Privacy Workshops (SPW) 2022; 80–87.

55. Abdallah M et al. TASHAROK: Using Mechanism Design for Enhancing Security Resource Allocation in Interdependent Systems. 2022 IEEE Symposium on Security and Privacy (SP) 2022; 249–266.

56. Sarvesh A, Carroll A, Gopalswamy S. Reshaping Local Path Planner. IEEE Robotics and Automation Letters 2022; 7(3):6534–6541.

57. Karnan H, Warnell G, Xiao X, Stone P. VOILA: Visual-Observation-Only Imitation Learning for Autonomous Navigation. 2022 International Conference on Robotics and Automation (ICRA) 2022; 2497–2503.

58. Triest S et al. TartanDrive: A Large-Scale Dataset for Learning Off-Road Dynamics Models. 2022 International Conference on Robotics and Automation (ICRA) 2022; 2546–2552.

59. Liang J et al. Search-Based Task Planning with Learned Skill Effect Models for Lifelong Robotic Manipulation. 2022 International Conference on Robotics and Automation (ICRA) 2022; 6351–6357.

60. Sikand K S et al. Visual Representation Learning for Preference-Aware Path Planning. 2022 International Conference on Robotics and Automation (ICRA) 2022; 11303–11309.

61. Vats S, Kroemer O, Likhachev . Synergistic Scheduling of Learning and Allocation of Tasks in Human-Robot Teams. 2022 International Conference on Robotics and Automation (ICRA) 2022; 2789–2795.

62. Qi X, Xie T, Mahloujifar S, Mittal P. Circumventing Backdoor Defenses That Are Based on Latent Separability. arXiv 2022; 2205.13613.

63. Qi X, Xie T, Mahloujifar S, Mittal P. Towards A Proactive ML Approach for Detecting Backdoor Poison Samples. arXiv 2022; 2205.13616.

64. Chen H, Zhang K, Hu S, You S, Kuo C J. Fake Satellite Image Detection via Parallel Subspace Learning (PSL). 2022 IEEE International Symposium on Circuits and Systems (ISCAS) 2022; 1:1502–1506.

65. Xu R et al. Smartadapt: Multi-branch Object Detection Framework for Videos on Mobiles. 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) 2022; 2518–2528.

66. Berger Z, Agrawal P, Liu T Y, Soatto S, Wong A. Stereoscopic Universal Perturbations across Different Architectures and Datasets. 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) 2022; 15159–15169.

67. Wang Z et al. Synthetic Generation of Face Videos with Plethysmograph Physiology. 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) 2022; 20555–20564.

68. Handelman D A et al. Adaptive human-robot teaming through integrated symbolic and subsymbolic artificial intelligence: preliminary results. Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications IV 2022; 13.

69. LaGrassa A, Kroemer O. Learning Model Preconditions for Planning with Multiple Models. arXiv 2022; 2206.05573.

70. Zhang Z et al. Neurotoxin: Durable Backdoors in Federated Learning. arXiv 2022; 2206.10341.

71. Cianfarani C, Bhagoji A N, Sehwag V, Zhao B Y, Mittal P. Understanding Robust Learning through the Lens of Representation Similarities. arXiv 2022; 2206.09868.

72. Liang J, Cheng X, Kroemer O. Learning Preconditions of Hybrid Force-Velocity Controllers for Contact-Rich Manipulation. arXiv 2022; 2206.12728.

73. Sommer D, Song L, Wagh S, Mittal P. Athena: Probabilistic Verification of Machine Unlearning. Proceedings on Privacy Enhancing Technologies 2022; 2022(3):268–290.

74. Elkady M, Lim J, Inouye D I. Discrete Tree Flows via Tree-Structured Permutations. arXiv 2022; 2207.01744.

75. Cho W, Gong Z, Inouye D I. Cooperative Distribution Alignment via JSD Upper Bound. arXiv 2022; 2207.02286.

76. Das M P, Conover D M, Eum S, Kwon H, Likhachev M. MA3: Model-Accuracy Aware Anytime Planning with Simulation Verification for Navigating Complex Terrains. Proceedings of the International Symposium on Combinatorial Search 2022; 15(1):65–73.

77. Mukherjee S, Aine S, Likhachev M. ePA*SE: Edge-Based Parallel A* for Slow Evaluations. Proceedings of the International Symposium on Combinatorial Search 2022; 15(1):136–144.

78. Lin R, Zhou Z, You S, Rao R, Kuo C J. Geometrical Interpretation and Design of Multilayer Perceptrons. IEEE Transactions on Neural Networks and Learning Systems 2022; 35(2):2545–2559.

79. Gasparino M V et al. WayFAST: Navigation With Predictive Traversability in the Field. IEEE Robotics and Automation Letters 2022; 7(4):10651–10658.

80. Chour K et al. A Reactive Energy-Aware Rendezvous Planning Approach for Multi-Vehicle Teams. 2022 IEEE 18th International Conference on Automation Science and Engineering (CASE) 2022; 537–542.

81. Tang X et al. Machine Learning with Differentially Private Labels: Mechanisms and Frameworks. Proceedings on Privacy Enhancing Technologies 2022; 2022(4):332–350.

82. Li X et al. DAG-based Task Orchestration for Edge Computing. 2022 41st International Symposium on Reliable Distributed Systems (SRDS) 2022; 23–34.

83. Jalali B, Kadambi A, Roychowdhury V. Physics-AI symbiosis. Machine Learning: Science and Technology 2022; 3(4):041001.

84. Xiao X et al. Learning Model Predictive Controllers with Real-Time Attention for Real-World Navigation. arXiv 2022; 2209.10780.

85. Zhou Z, Fu H, You S, Kuo C J. GUSOT: Green and Unsupervised Single Object Tracking for Long Video Sequences. 2022 IEEE 24th International Workshop on Multimedia Signal Processing (MMSP) 2022; 25:1–6.

86. Lee J et al. Virtuoso: Energy- and Latency-aware Streamlining of Streaming Videos on Systems-on-Chips. ACM Transactions on Design Automation of Electronic Systems 2022; 28(3):1–32.

87. Kumar S, Iftekhar A, Prashnani E, Manjunath B. LOCL: Learning Object-Attribute Composition using Localization. arXiv 2022; 2210.03780.

88. Kulinski S, Inouye D I. Towards Explaining Distribution Shifts. arXiv 2022; 2210.10275.

89. Sharma M, Kroemer O. Efficiently Learning Manipulations by Selecting Structured Skill Representations. 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) 2022; 22:1039–1046.

90. Karnan H et al. VI-IKD: High-Speed Accurate Off-Road Navigation using Learned Visual-Inertial Inverse Kinodynamics. 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) 2022; 3294–3301.

91. Li T, Krakow L W, Gopalswamy S. SMA-NBO: A Sequential Multi-Agent Planning with Nominal Belief-State Optimization in Target Tracking. 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) 2022.

92. Atreya P et al. High-Speed Accurate Robot Control using Learned Forward Kinodynamics and Non-linear Least Squares Optimization. 2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) 2022; 11789–11795.

93. Wu T, Wang T, Sehwag V, Mahloujifar S, Mittal P. Just Rotate it: Deploying Backdoor Attacks via Rotation Transformation. Proceedings of the 15th ACM Workshop on Artificial Intelligence and Security 2022; 91–102.

94. Wang S, Hosseinalipour S, Gorlatova M, Brinton C G, Chiang M. UAV-Assisted Online Machine Learning Over Multi-Tiered Networks: A Hierarchical Nested Personalized Federated Learning Approach. IEEE Transactions on Network and Service Management 2022; 20(2):1847–1865.

95. Fu H et al. Acceleration of Subspace Learning Machine via Particle Swarm Optimization and Parallel Processing. 2022 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC) 2022; 1019–1024.

96. Nair A et al. DynaBARN: Benchmarking Metric Ground Navigation in Dynamic Environments. 2022 IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR) 2022; 347–352.

97. Wagle S, Hosseinalipour S, Khosravan N, Chiang M, Brinton C G. Embedding Alignment for Unsupervised Federated Learning via Smart Data Exchange. GLOBECOM 2022 - 2022 IEEE Global Communications Conference 2022.

98. Yu E CG, Deng X, Su Y. Don't Generate, Discriminate: A Proposal for Grounding Language Models to Real-World Environments. arXiv 2022; 2212.09736.

99. Chari P et al. On Learning Mechanical Laws of Motion From Video Using Neural Networks. IEEE Access 2023; 11:30129–30145.

100. Shah K, Shah A, Lau C P, Melo C M D, Chellappa R. Multi-View Action Recognition using Contrastive Learning. 2023 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) 2023; 3370–3380.

101. Shen Y et al. Archangel: A Hybrid UAV-Based Human Detection Benchmark With Position and Pose Metadata. IEEE Access 2023; 11:80958–80972.

102. Narayanan P et al. A Multi-Purpose Realistic Haze Benchmark With Quantifiable Haze Levels and Ground Truth. IEEE Transactions on Image Processing 2023; 32:3481–3492.

103. Wijk D V, Eves K, Valasek J. Deep Reinforcement Learning Controller for Autonomous Tracking of Evasive Ground Target. AIAA SCITECH 2023 Forum 2023.

104. Wu T et al. Uncovering Adversarial Risks of Test-Time Adaptation. arXiv 2023; 2301.12576.

105. Debenedetti E, Sehwag V, Mittal P. A Light Recipe to Train Robust Vision Transformers. 2023 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML) 2023; 225–253.

106. Dai S et al. Characterizing the Optimal 0-1 Loss for Multi-class Classification with a Test-time Attacker. arXiv 2023; 2302.10722.

107. Dai S et al. MultiRobustBench: Benchmarking Robustness Against Multiple Attacks. arXiv 2023; 2302.10980.

108. Jeyakumar J V, Sarker A, Garca L A R, Srivastava M. X-CHAR. Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies 2023; 7(1):1–28.

109. Wang J T, Mahloujifar S, Wu T, Jia R, Mittal P. A Randomized Approach for Tight Privacy Accounting. arXiv 2023; 2304.07927.

110. Ganguly B et al. Multi-Edge Server-Assisted Dynamic Federated Learning With an Optimized Floating Aggregation Point. IEEE/ACM Transactions on Networking 2023; 31(6):2682–2697.

111. Katyal K D et al. Leveraging synthetic data for robust gesture recognition. Synthetic Data for Artificial Intelligence and Machine Learning: Tools, Techniques, and Applications 2023; 37.

112. Xiang C, Valtchanov A, Mahloujifar S, Mittal P. ObjectSeeker: Certifiably Robust Object Detection against Patch Hiding Attacks via Patch-agnostic Masking. 2023 IEEE Symposium on Security and Privacy (SP) 2023; 1329–1347.

113. Lin Y, Prasanna V K. HyScale-GNN: A Scalable Hybrid GNN Training System on Single-Node Heterogeneous Architecture. 2023 IEEE International Parallel and Distributed Processing Symposium (IPDPS) 2023; 557–567.

114. Panda A, Wu T, Wang J T, Mittal P. Privacy-Preserving In-Context Learning for Large Language Models. arXiv 2023; 2305.01639.

115. Chour K et al. An agent-based modeling framework for the multi-UAV rendezvous recharging problem. Robotics and Autonomous Systems 2023; 166:104442.

116. Fu H, Yang Y, Mishra V K, Kuo C J. Classification via Subspace Learning Machine (SLM): Methodology and Performance Evaluation. ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) 2023; 1:1–5.

117. Wang S et al. Toward Cooperative Federated Learning Over Heterogeneous Edge/Fog Networks. IEEE Communications Magazine 2023; 61(12):54–60.

118. Lim H et al. An Empirical Study of 5G: Effect of Edge on Transport Protocol and Application Performance. IEEE Transactions on Mobile Computing 2023; 23(4):3172–3186.

119. Zhou H, Kannan R, Swami A, Prasanna V. HTNet: Dynamic WLAN Performance Prediction using Heterogenous Temporal GNN. IEEE INFOCOM 2023 - IEEE Conference on Computer Communications 2023; 1–10.

120. Mitrano P, LaGrassa A, Kroemer O, Berenson D. Focused Adaptation of Dynamics Models for Deformable Object Manipulation. 2023 IEEE International Conference on Robotics and Automation (ICRA) 2023; 5931–5937.

121. Xu Z, Liu B, Xiao X, Nair A, Stone P. Benchmarking Reinforcement Learning Techniques for Autonomous Navigation. 2023 IEEE International Conference on Robotics and Automation (ICRA) 2023; 9224–9230.

122. Castro M G et al. How Does It Feel? Self-Supervised Costmap Learning for Off-Road Vehicle Traversability. 2023 IEEE International Conference on Robotics and Automation (ICRA) 2023; 931–938.

123. Reddy A et al. Synthetic-to-Real Domain Adaptation for Action Recognition: A Dataset and Baseline Performances. 2023 IEEE International Conference on Robotics and Automation (ICRA) 2023; 11374–11381.

124. Singh I et al. ProgPrompt: Generating Situated Robot Task Plans using Large Language Models. 2023 IEEE International Conference on Robotics and Automation (ICRA) 2023; 11523–11530.

125. Vats S, Likhachev M, Kroemer O. Efficient Recovery Learning using Model Predictive Meta-Reasoning. 2023 IEEE International Conference on Robotics and Automation (ICRA) 2023; 12:7258–7264.

126. Triest S et al. Learning Risk-Aware Costmaps via Inverse Reinforcement Learning for Off-Road Navigation. 2023 IEEE International Conference on Robotics and Automation (ICRA) 2023; 924–930.

127. Natarajan R, Johnston G L H, Simaan N, Likhachev M, Choset H. Torque-Limited Manipulation Planning through Contact by Interleaving Graph Search and Trajectory Optimization. 2023 IEEE International Conference on Robotics and Automation (ICRA) 2023; 8148–8154.

128. Saxena D M, Likhachev M. Planning for Complex Non-prehensile Manipulation Among Movable Objects by Interleaving Multi-Agent Pathfinding and Physics-Based Simulation. 2023 IEEE International Conference on Robotics and Automation (ICRA) 2023; 8141–8147.

129. Khajenejad M, Brown S, Martínez S. Distributed Interval Observers for Bounded-Error LTI Systems. 2023 American Control Conference (ACC) 2023; 121:2444–2449.

130. Khajenejad M, Brown S, Martínez S. Distributed Resilient Interval Observers for Bounded-Error LTI Systems Subject to False Data Injection Attacks. 2023 American Control Conference (ACC) 2023; 3502–3507.

131. Kulinski S, Waytowich N R, Hare J Z, Inouye D I. StarCraftImage: A Dataset For Prototyping Spatial Reasoning Methods For Multi-Agent Environments. 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) 2023; 22004–22013.

132. Kadambi A, Melo C P D, Hsieh C, Srivastava M, Soatto S. Incorporating physics into data-driven computer vision. Nature Machine Intelligence 2023; 5(6):572–580.

133. Zhang H et al. WeatherStream: Light Transport Automation of Single Image Deweathering. 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) 2023; 13499–13509.

134. Zhao J C et al. The Resource Problem of Using Linear Layer Leakage Attack in Federated Learning. 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) 2023; 3974–3983.

135. Kumar S, Arevalo I, Iftekhar A, Manjunath B S. MethaneMapper: Spectral Absorption Aware Hyperspectral Transformer for Methane Detection. 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) 2023; 17609–17618.

136. Wang Z et al. ALTO: Alternating Latent Topologies for Implicit 3D Reconstruction. 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) 2023; 259–270.

137. Bandarupalli A, Jain S, Melachuri A, Pappas J, Chaterji S. Vega: Drone-based Multi-Altitude Target Detection for Autonomous Surveillance. 2023 19th International Conference on Distributed Computing in Smart Systems and the Internet of Things (DCOSS-IoT) 2023.

138. Deng X et al. Mind2Web: Towards a Generalist Agent for the Web. arXiv 2023; 2306.06070.

139. Handelman D A et al. Multi-agent playbook for human-robot teaming. Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications V 2023; 14.

140. Iftekhar A, Kumar S, McEver R A, You S, Manjunath B S. Gtnet: guided transformer network for detecting human-object interactions. Pattern Recognition and Tracking XXXIV 2023; 24.

141. Abraham S J et al. Efficient hyperparameter optimization for ATR using homotopy parametrization. Automatic Target Recognition XXXIII 2023; 6.

142. Zhang K, Mo L, Chen W, Sun H, Su Y. MagicBrush: A Manually Annotated Dataset for Instruction-Guided Image Editing. arXiv 2023; 2306.10012.

143. Kulinski S, Zhou Z, Bai R, Kocaoğlu M, Inouye D I. Towards Characterizing Domain Counterfactuals For Invertible Latent Causal Models. arXiv 2023; 2306.11281.

144. Suresh A, Nieto C, Martínez S. PuRR-RRT*: Behavioral Path Planning in Uncertain, Risky and Rewarding Environments. 2023 20th International Conference on Ubiquitous Robots (UR) 2023; 211–216.

145. Suresh A, Taylor A, Riek L D, Martínez S. Robot Navigation in Risky, Crowded Environments: Understanding Human Preferences. IEEE Robotics and Automation Letters 2023; 8(9):5632–5639.

146. Saxena D, Likhachev M. Planning for Manipulation among Movable Objects: Deciding Which Objects Go Where, in What Order, and How. Proceedings of the International Conference on Automated Planning and Scheduling 2023; 33(1):668–676.

147. Yang H, Mukherjee S, Likhachev M. A-ePA*SE: Anytime Edge-Based Parallel A* for Slow Evaluations. Proceedings of the International Symposium on Combinatorial Search 2023; 16(1):163–167.

148. Mukherjee S, Likhachev M. GePA*SE: Generalized Edge-Based Parallel A* for Slow Evaluations. Proceedings of the International Symposium on Combinatorial Search 2023; 16(1):153–157.

149. Sharma A et al. FLAIR: Defense against Model Poisoning Attack in Federated Learning. Proceedings of the ACM Asia Conference on Computer and Communications Security 2023; 553–566.

150. Hosseinalipour S et al. Parallel Successive Learning for Dynamic Distributed Model Training Over Heterogeneous Wireless Networks. IEEE/ACM Transactions on Networking 2023; 32(1):222–237.

151. Bai R, Bagchi S, Inouye D I. Benchmarking Algorithms for Federated Domain Generalization. arXiv 2023; 2307.04942.

152. Xie T et al. BaDExpert: Extracting Backdoor Functionality for Accurate Backdoor Input Detection. arXiv 2023; 2308.12439.

153. Lukin S M et al. Navigating to Success in Multi-Modal Human-Robot Collaboration: Analysis and Corpus Release. 2023 32nd IEEE International Conference on Robot and Human Interactive Communication (RO-MAN) 2023; 1859–1865.

154. Wang J T, Zhu Y, Wang Y, Jia R, Mittal P. Threshold KNN-Shapley: A Linear-Time and Privacy-Friendly Approach to Data Valuation. arXiv 2023; 2308.15709.

155. Zhou Z, You S, Kuo C J. Unsupervised Green Object Tracker (GOT) without Offline Pre-training. arXiv 2023; 2309.09078.

156. Zhang B, Kannan R, Prasanna V K, Busart C. Accelerating GNN-Based SAR Automatic Target Recognition on HBM-Enabled FPGA. 2023 IEEE High Performance Extreme Computing Conference (HPEC) 2023; 1–7.

157. Parikh D, Zhang B, Kannan R, Prasanna V K, Busart C. Performance of Graph Neural Networks for Point Cloud Applications. 2023 IEEE High Performance Extreme Computing Conference (HPEC) 2023; 1–7.

158. Gupta N, Zhang P, Kannan R, Prasanna V K. PaCKD: Pattern-Clustered Knowledge Distillation for Compressing Memory Access Prediction Models. 2023 IEEE High Performance Extreme Computing Conference (HPEC) 2023; 1–7.

159. Wiggins S, Meng Y, Kannan R, Prasanna V K. Accelerating Multi-Agent DDPG on CPU-FPGA Heterogeneous Platform. 2023 IEEE High Performance Extreme Computing Conference (HPEC) 2023; 1–7.

160. Fein-Ashley J, Tian Y, Kannan R, Prasanna V K, Busart C. Benchmarking Deep Learning Classifiers for SAR Automatic Target Recognition. 2023 IEEE High Performance Extreme Computing Conference (HPEC) 2023; 1–6.

161. Gorle A R, Zhang P, Kannan R, Prasanna V K. G-MAP: A Graph Neural Network-Based Framework for Memory Access Prediction. 2023 IEEE High Performance Extreme Computing Conference (HPEC) 2023; abs/1803.02329:1–7.

162. Karnan H et al. STERLING: Self-Supervised Terrain Representation Learning from Unconstrained Robot Experience. arXiv 2023; 2309.15302.

163. Parasnis R, Hosseinalipour S, Chu Y, Chiang M, Brinton C G. Connectivity-Aware Semi-Decentralized Federated Learning over Time-Varying D2D Networks. Proceedings of the Twenty-fourth International Symposium on Theory, Algorithmic Foundations, and Protocol Design for Mobile Networks and Mobile Computing 2023; 31–40.

164. Gutierrez F et al. Learned Compressive Representations for Single-Photon 3D Imaging. 2023 IEEE/CVF International Conference on Computer Vision (ICCV) 2023; 10722–10732.

165. Phatak R, Shell D A. Sensor Selection for Fine-Grained Behavior Verification that Respects Privacy. 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) 2023; 88:10117–10124.

166. Song C H et al. LLM-Planner: Few-Shot Grounded Planning for Embodied Agents with Large Language Models. 2023 IEEE/CVF International Conference on Computer Vision (ICCV) 2023; 2986–2997.

167. Dutson M, Li Y, Gupta M. Eventful Transformers: Leveraging Temporal Redundancy in Vision Transformers. 2023 IEEE/CVF International Conference on Computer Vision (ICCV) 2023; 16865–16877.

168. Li T, Krakow L W, Gopalswamy S. Emergent Cooperative Behavior in Distributed Target Tracking with Unknown Occlusions. 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) 2023; 13:665–672.

169. Abraham S J, Maduranga K D G, Kinnison J, Hauenstein J D, Scheirer W J. NCQS: Nonlinear Convex Quadrature Surrogate Hyperparameter Optimization. 2023 IEEE/CVF International Conference on Computer Vision Workshops (ICCVW) 2023; 24:1187–1195.

170. Qi X et al. Fine-tuning Aligned Language Models Compromises Safety, Even When Users Do Not Intend To!. arXiv 2023; 2310.03693.

171. Gilani S N, Pollard K A, Traum D. Multimodal Prediction of User's Performance in High-Stress Dialogue Interactions. International Cconference on Multimodal Interaction 2023; 71–75.

172. Ruschel R, Iftekhar A, Manjunath B S, You S. BLoad: Enhancing Neural Network Training with Efficient Sequential Data Handling. arXiv 2023; 2310.10879.

173. Wijeratne S, Kannan R, Prasanna V K. Dynasor: A Dynamic Memory Layout for Accelerating Sparse MTTKRP for Tensor Decomposition on Multi-core

CPU. 2023 IEEE 35th International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD) 2023; 23–33.

174. Pahuja V, Wang B, Latapie H, Srinivasa J, Su Y. A Retrieve-and-Read Framework for Knowledge Graph Link Prediction. Proceedings of the 32nd ACM International Conference on Information and Knowledge Management 2023; 1992–2002.

175. Wang T, Kannan R, Prasanna V K. Training Heterogeneous Graph Neural Networks using Bandit Sampling. Proceedings of the 32nd ACM International Conference on Information and Knowledge Management 2023; 4345–4349.

176. Medellin A, Bhamri A, Langari R, Gopalswamy S. Real-Time Semantic Segmentation Using Hyperspectral Images for Unstructured and Unknown Environments. 2023 13th Workshop on Hyperspectral Imaging and Signal Processing: Evolution in Remote Sensing (WHISPERS) 2023; 1–5.

177. Maheshwari P et al. PIAug – Physics Informed Augmentation for Learning Vehicle Dynamics for Off-Road Navigation. arXiv 2023; 2311.00815.

178. Ye T, Kannan R, Prasanna V K, Busart C, Kaplan L. Realistic Scatterer Based Adversarial Attacks on SAR Image Classifiers. 2023 IEEE International Radar Conference (RADAR) 2023; 1–6.

179. Zhang B et al. How can Human-in-the-loop Improve the Performance of SAR ATR? A Reinforcement Learning Based Approach. 2023 IEEE International Radar Conference (RADAR) 2023; 1–6.

180. Deng G et al. An Efficient Distributed Graph Engine for Deep Learning on Graphs. Proceedings of the SC ’23 Workshops of the International Conference on High Performance Computing, Network, Storage, and Analysis 2023; 922–931.

181. Han D et al. Federated Split Learning With Joint Personalization-Generalization for Inference-Stage Optimization in Wireless Edge Networks. IEEE Transactions on Mobile Computing 2023; 23(6):7048–7065.

182. Baroni P, Cerutti F, Giacomin M, Kaplan L, Şensoy M. A Formal Account for Reasoning About Uncertain Probabilities in Intelligence Analysis. 2023 IEEE

International Workshop on Technologies for Defense and Security (TechDefense) 2023; 434–438.

183. Cominelli M et al. PRIVAtE: Passive Radar Interpretability using Variational Auto Encoders. 2023 IEEE International Workshop on Technologies for Defense and Security (TechDefense) 2023; 444–448.

184. Nodari L, Cerutti F. Assessing the Robustness of Intelligence-Driven Reinforcement Learning. 2023 IEEE International Workshop on Technologies for Defense and Security (TechDefense) 2023; 32:439–443.

185. Bhat A et al. EESMR. Proceedings of the 24th International Middleware Conference on ZZZ 2023; 1–14.

186. Bhat A, Bandarupalli A, Bagchi S, Kate A, Reiter M K. The Unique Chain Rule andIts Applications. Lecture Notes in Computer Science 2023; 38–55.

187. Xiao X et al. Autonomous Ground Navigation in Highly Constrained Spaces: Lessons Learned From the Second BARN Challenge at ICRA 2023 [Competitions]. IEEE Robotics and Automation Magazine 2023; 30(4):91–97.

188. Fu H, Wang X, Mishra V K, Kuo C J. Subspace Learning Machine with Soft Partitioning (SLM/SP): Methodology and Performance Benchmarking. 2023 IEEE International Conference on Visual Communications and Image Processing (VCIP) 2023; 1–5.

189. Bhamri A et al. Unifying Ontological Framework for All-Terrain Datasets. 2023 Seventh IEEE International Conference on Robotic Computing (IRC) 2023; 292–296.

190. Li J et al. Scenario-Game ADMM: A Parallelized Scenario-Based Solver for Stochastic Noncooperative Games. 2023 62nd IEEE Conference on Decision and Control (CDC) 2023; 8093–8099.

191. Chen S, Yu Y, Fridovich-Keil D, Topcu U. Soft-Bellman Equilibrium in Affine Markov Games: Forward Solutions and Inverse Learning. 2023 62nd IEEE Conference on Decision and Control (CDC) 2023; 2202–2207.

192. Wang X, Mishra V K, Kuo C J. Enhancing Edge Intelligence with Highly Discriminant LNT Features. 2023 IEEE International Conference on Big Data (BigData) 2023; 3880–3887.

193. Gao P F, Jaafar A, Reily B, Reardon C, Zhang H. Compositional Zero-Shot Learning for Attribute-Based Object Reference in Human-Robot Interaction. arXiv 2023; 2312.13655.

194. Nguemkam A L, Hemida A, Tchendji V K, Tosh D K, Kamhoua C. Optimal Honeypot Allocation Using Core Attack Graph in Partially Observable Stochastic Games. IEEE Access 2024; 12:187444–187455.

195. Sun Y, Liu T, Kumar P R, Shahrampour S. Linear Convergence of Independent Natural Policy Gradient in Games With Entropy Regularization. IEEE Control Systems Letters 2024; 8:1217–1222.

196. Ardon L, Furelos-Blanco D, Russo A. Learning Reward Machines in Cooperative Multi-agent Tasks. Lecture Notes in Computer Science 2024; 43–59.

197. Mvah F et al. Countering ARP spoofing attacks in software-defined networks using a game-theoretic approach. Computers and Security 2024; 139:103696.

198. Lee J J, Rathinam S. A Meta-Heuristic Approach for an Aerial-Ground Vehicle Path Planning Problem. AIAA SCITECH 2024 Forum 2024.

199. Fu H, Yang Y, Mishra V K, Kuo C J. Subspace learning machine (SLM): Methodology and performance evaluation. Journal of Visual Communication and Image Representation 2024; 98:104058.

200. Saxena S, Sharma M, Kroemer O. MResT: Multi-Resolution Sensing for Real-Time Control with Vision-Language Models. arXiv 2024; 2401.14502.

201. Thakkar R S, Samyal A S, Fridovich-Keil D, Xu Z, Topcu U. Hierarchical Control for Cooperative Teams in Competitive Autonomous Racing. IEEE Transactions on Intelligent Vehicles 2024; 9(5):4845–4860.

202. Han D, Hosseinalipour S, Love D J, Chiang M, Brinton C G. Cooperative Federated Learning Over Ground-to-Satellite Integrated Networks: Joint Local Computation and Data Offloading. IEEE Journal on Selected Areas in Communications 2024; 42(5):1080–1096.

203. Mvah F et al. Deception-Based IDS Against ARP Spoofing Attacks in Software-Defined Networks. 2024 International Conference on Computing, Networking and Communications (ICNC) 2024; 188–192.

204. Xu Z et al. Towards Few-Shot Adaptation of Foundation Models via Multitask Finetuning. arXiv 2024; 2402.15017.

205. Panda A, Choquette-Choo C A, Zhang Z, Yang Y, Mittal P. Teach LLMs to Phish: Stealing Private Information from Language Models. arXiv 2024; 2403.00871.

206. Parasnis R, Hosseinalipour S, Chu Y, Chiang M, Brinton C G. Energy-Efficient Connectivity-Aware Learning Over Time-Varying D2D Networks. IEEE Journal of Selected Topics in Signal Processing 2024; 18(2):242–258.

207. Areti M, Hasnain Z. Leveraging local neighborhood features in 3D unstructured point cloud data for geometry-based automated damage delineation. Automation in Construction 2024; 162:105393.

208. López A M, Fridovich-Keil D. Decomposing Control Lyapunov Functions for Efficient Reinforcement Learning. arXiv 2024; 2403.12210.

209. Qi X et al. Visual Adversarial Examples Jailbreak Aligned Large Language Models. Proceedings of the AAAI Conference on Artificial Intelligence 2024; 38(19):21527–21536.

210. Zhang B, Kannan R, Busart C, Prasanna V. GCV-Turbo: End-to-end Acceleration of GNN-based Computer Vision Tasks on FPGA. 2024 IEEE 32nd Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM) 2024; 66–77.

211. Parikh D et al. Accelerating ViT Inference on FPGA through Static and Dynamic Pruning. 2024 IEEE 32nd Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM) 2024; 78–89.

212. Wang X, Ye T, Kannan R, Prasanna V. FACTUAL: A Novel Framework for Contrastive Learning Based Robust SAR Image Classification. 2024 IEEE Radar Conference (RadarConf24) 2024; 1–6.

213. Ye T, Kannan R, Prasanna V, Busart C. Uncertainty-Aware SAR ATR: Defending Against Adversarial Attacks via Bayesian Neural Networks. 2024 IEEE Radar Conference (RadarConf24) 2024; 1–6.

214. Lin Y, Deng G, Prasanna V K. A Unified CPU-GPU Protocol for GNN Training. Proceedings of the 21st ACM International Conference on Computing Frontiers 2024; 155–163.

215. Gupta N, Kannan N, Zhang P, Prasanna V K. TabConv: Low-Computation CNN Inference via Table Lookups. Proceedings of the 21st ACM International Conference on Computing Frontiers 2024; 180–188.

216. Wijeratne S, Kannan R, Prasanna V K. Sparse MTTKRP Acceleration for Tensor Decomposition on GPU. Proceedings of the 21st ACM International Conference on Computing Frontiers 2024; 88–96.

217. Limbu M, Hu Z, Wang X, Shishika D, Xiao X. Scaling Team Coordination on Graphs with Reinforcement Learning. 2024 IEEE International Conference on Robotics and Automation (ICRA) 2024; 16538–16544.

218. Liang J et al. MTG: Mapless Trajectory Generator with Traversability Coverage for Outdoor Navigation. 2024 IEEE International Conference on Robotics and Automation (ICRA) 2024; 2396–2402.

219. Karnan H, Yang E, Warnell G, Biswas J, Stone P. Wait, That Feels Familiar: Learning to Extrapolate Human Preferences for Preference-Aligned Path Planning. 2024 IEEE International Conference on Robotics and Automation (ICRA) 2024; 13008–13014.

220. Gu Q et al. ConceptGraphs: Open-Vocabulary 3D Scene Graphs for Perception and Planning. 2024 IEEE International Conference on Robotics and Automation (ICRA) 2024; 5021–5028.

221. Datar A, Pan C, Nazeri M, Xiao X. Toward Wheeled Mobility on Vertically Challenging Terrain: Platforms, Datasets, and Algorithms. 2024 IEEE International Conference on Robotics and Automation (ICRA) 2024; 16322–16329.

222. Xu Z, Raj A H, Xiao X, Stone P. Dexterous Legged Locomotion in Confined 3D Spaces with Reinforcement Learning. 2024 IEEE International Conference on Robotics and Automation (ICRA) 2024; 11474–11480.

223. LaGrassa A, Lee M Y, Kroemer O. Task-Oriented Active Learning of Model Preconditions for Inaccurate Dynamics Models. 2024 IEEE International Conference on Robotics and Automation (ICRA) 2024; 16445.

224. Saxena D M, Likhachev M. Improved M4M: Faster and Richer Planning for Manipulation Among Movable Objects in Cluttered 3D Workspaces. 2024 IEEE International Conference on Robotics and Automation (ICRA) 2024; 87:10902–10909.

225. Sivaprakasam M et al. TartanDrive 2.0: More Modalities and Better Infrastructure to Further Self-Supervised Learning Research in Off-Road Driving Tasks. 2024 IEEE International Conference on Robotics and Automation (ICRA) 2024; 12606.

226. Das D, Lu Y, Plaku E, Xiao X. Motion Memory: Leveraging Past Experiences to Accelerate Future Motion Planning. 2024 IEEE International Conference on Robotics and Automation (ICRA) 2024; 16467–16474.

227. Noseworthy M D, Shaw S, Kessens C C, Roy N. Amortized Inference for Efficient Grasp Model Adaptation. 2024 IEEE International Conference on Robotics and Automation (ICRA) 2024; 1886–1892.

228. Seker M Y, Kroemer O. Estimating Material Properties of Interacting Objects Using Sum-GP-UCB. 2024 IEEE International Conference on Robotics and Automation (ICRA) 2024; 7:16684–16690.

229. Raj A H et al. Rethinking Social Robot Navigation: Leveraging the Best of Two Worlds. 2024 IEEE International Conference on Robotics and Automation (ICRA) 2024; 16330–16337.

230. Zhao J C et al. Loki: Large-scale Data Reconstruction Attack against Federated Learning through Model Manipulation. 2024 IEEE Symposium on Security and Privacy (SP) 2024; 1287–1305.

231. Deng G et al. TASER: Temporal Adaptive Sampling for Fast and Accurate Dynamic Graph Representation Learning. 2024 IEEE International Parallel and Distributed Processing Symposium (IPDPS) 2024; 926–937.

232. Bandara W G C, Patel V M. Attention Prompt Tuning: Parameter-efficient Adaptation of Pre-trained Models for Action Recognition. 2024 IEEE 18th

International Conference on Automatic Face and Gesture Recognition (FG) 2024; 1–10.

233. Zhang P, Gupta N, Kannan R, Prasanna V K. Attention, Distillation, and Tabularization: Towards Practical Neural Network-Based Prefetching. 2024 IEEE International Parallel and Distributed Processing Symposium (IPDPS) 2024; 876–888.

234. Pappas J et al. Adaptive object detection algorithms for resource constrained autonomous robotic systems. Disruptive Technologies in Information Sciences VIII 2024; 28:10.

235. Handelman D A et al. Leveraging foundation models for scene understanding in human-robot teaming. Artificial Intelligence and Machine Learning for Multi-Domain Operations Applications VI 2024; 32.

236. Reddy A et al. An evaluation of large pre-trained models for gesture recognition using synthetic videos. Synthetic Data for Artificial Intelligence and Machine Learning: Tools, Techniques, and Applications II 2024; 16.

237. Mu F et al. Towards 3D Vision with Low-Cost Single-Photon Cameras. 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) 2024; 5302–5311.

238. Zhang C, Stepputtis S, Campbell J, Sycara K, Xie Y. HiKER-SGG: Hierarchical Knowledge Enhanced Robust Scene Graph Generation. 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) 2024; 28233–28243.

239. Kil J et al. Dual-View Visual Contextualization for Web Navigation. 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) 2024; 14445–14454.

240. Zhao J C, Dabholkar A, Sharma A, Bagchi S. Leak and Learn: An Attacker's Cookbook to Train Using Leaked Data from Federated Learning. 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) 2024; 33:12247–12256.

241. Reddy A et al. Unsupervised Video Domain Adaptation with Masked Pre-Training and Collaborative Self-Training. 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) 2024; 18919–18929.

242. Zhou S et al. Feature 3DGS: Supercharging 3D Gaussian Splatting to Enable Distilled Feature Fields. 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) 2024; 21676–21685.

243. Ranasinghe Y, Hegde D, Patel V M. MonoDiff: Monocular 3D Object Detection and Pose Estimation with Diffusion Models. 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) 2024; 10659–10670.

244. Mu F, Mo S, Li Y. SnAG: Scalable and Accurate Video Grounding. 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) 2024; 18930–18940.

245. Cominelli M et al. Neuro-Symbolic Fusion of Wi-Fi Sensing Data for Passive Radar with Inter-Modal Knowledge Transfer. 2024 27th International Conference on Information Fusion (FUSION) 2024; 1–8.

246. Bandarupalli A et al. SensorBFT: Fault-Tolerant Target Localization Using Voronoi Diagrams and Approximate Agreement. 2024 IEEE 44th International Conference on Distributed Computing Systems (ICDCS) 2024; 186–197.

247. Suresh A, Nieto-Granda C, Martínez S. Robotic Exploration Using Generalized Behavioral Entropy. IEEE Robotics and Automation Letters 2024; 9(9):8011–8018.

248. Abdallah M et al. Game Theory in Distributed Systems Security: Foundations, Challenges, and Future Directions. IEEE Security and Privacy 2024; 23(1):64–74.

249. Xiao X et al. Autonomous Ground Navigation in Highly Constrained Spaces: Lessons Learned From the Third BARN Challenge at ICRA 2024 [Competitions]. IEEE Robotics and Automation Magazine 2024; 31(3):197–204.

250. Dabholkar A et al. Adversarial attacks on reinforcement learning agents for command and control. The Journal of Defense Modeling and Simulation: Applications, Methodology, Technology 2024.

251. Ganji S, Kumar P R. Seeing the Unseen: The REVEAL Protocol to Expose the Wireless Man-in-the-Middle. IEEE Transactions on Wireless Communications 2024; 23(11):17143–17156.

252. Pokhrel A, Nazeri M, Datar A, Xiao X. CAHSOR: Competence-Aware High-Speed Off-Road Ground Navigation in SE(3). IEEE Robotics and Automation Letters 2024; 9(11):9653–9660.

253. Zhang Z, Parikh D, Zhang Y, Prasanna V. Benchmarking the Performance of Large Language Models on the Cerebras Wafer Scale Engine. 2024 IEEE High Performance Extreme Computing Conference (HPEC) 2024; 1–7.

254. Jayanth R, Gupta N, Prasanna V K. Benchmarking Edge AI Platforms for High-Performance ML Inference. 2024 IEEE High Performance Extreme Computing Conference (HPEC) 2024; 1–7.

255. Zhang B, Kannan R, Busart C, Prasanna V K. VisionAGILE: A Versatile Domain-Specific Accelerator for Computer Vision Tasks. IEEE Transactions on Parallel and Distributed Systems 2024; 35(12):2405–2422.

256. Fein-Ashley J, Wickramasinghe S, Zhang B, Kannan R, Prasanna V K. A Single Graph Convolution is All You Need: Efficient Grayscale Image Classification. 2024 IEEE International Conference on Image Processing (ICIP) 2024; 849–855.

257. Zhou S et al. DreamScene360: Unconstrained Text-to-3D Scene Generation with Panoramic Gaussian Splatting. Lecture Notes in Computer Science 2024; 324–342.

258. Nazeri M, Wang J, Payandeh A, Xiao X. VANP: Learning Where to See for Navigation with Self-Supervised Vision-Action Pre-Training. 2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) 2024; 2741–2746.

259. Zhou Y et al. Team Coordination on Graphs: Problem, Analysis, and Algorithms. 2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) 2024; 5748–5755.

260. Datar A, Pan C, Nazeri M, Pokhrel A, Xiao X. Terrain-Attentive Learning for Efficient 6-DoF Kinodynamic Modeling on Vertically Challenging Terrain. 2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) 2024; 5438–5443.

261. Wijeratne S, Kannan R, Prasanna V K. Accelerating Sparse MTTKRP for Small Tensor Decomposition on GPU $\{}mathcal{X}_{(d)}$. 2024 58th Asilomar Conference on Signals, Systems, and Computers 2024; 1900–1905.

262. Barbin A, Cerutti F, Gerevini A. Learning Reliable PDDL Models for Classical Planning from Visual Data. 2024 IEEE 36th International Conference on Tools with Artificial Intelligence (ICTAI) 2024; 722–728.

263. Wang B et al. Adapting Complex Event Detection to Perceptual Domain Shifts. MILCOM 2024 - 2024 IEEE Military Communications Conference (MILCOM) 2024; 1–6.

264. Lukin S M et al. SCOUT: A Situated and Multi-Modal Human-Robot Dialogue Corpus. arXiv 2024; 2411.12844.

265. Lu C et al. RECON: Training-Free Acceleration for Text-to-Image Synthesis withRetrieval of Concept Prompt Trajectories. Lecture Notes in Computer Science 2024; 288–306.

266. Bandarupalli A, Bhat A, Bagchi S, Kate A, Reiter M K. Random Beacons in Monte Carlo: Efficient Asynchronous Random Beacon *without* Threshold Cryptography. Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security 2024; 2621–2635.

267. Song D et al. VLM-Social-Nav: Socially Aware Robot Navigation Through Scoring Using Vision-Language Models. IEEE Robotics and Automation Letters 2024; 10(1):508–515.

268. Şensoy M, Kaplan L, Julier S, Saleki M, Cerutti F. Risk-aware classification via uncertainty quantification. Expert Systems with Applications 2024; 265:125906.

269. Wang T, Kannan R, Prasanna V. Heterogeneous Graph Neural Network based on Bandit Sampling for Attribute Completion. 2024 IEEE International Conference on Knowledge Graph (ICKG) 2024; 404–411.

270. Lin T, Kumar P R. Dynamic Watermarking for Cyber-Security of Nonlinear Stochastic Systems. 2024 IEEE 63rd Conference on Decision and Control (CDC) 2024; 2572–2577.

271. Yu Y, Thorpe A J, Milzman J, Fridovich-Keil D, Topcu U. Sensing Resource Allocation Against Data-Poisoning Attacks in Traffic Routing. 2024 IEEE 63rd Conference on Decision and Control (CDC) 2024; 8396–8402.

272. Datar A, Pan C, Xiao X. Learning to Model and Plan for Wheeled Mobility on Vertically Challenging Terrain. IEEE Robotics and Automation Letters 2024; 10(2):1505–1512.

273. Zhou H, Zhang B, Kannan R, Prasanna V, Busart C. ViTeGNN: Towards Versatile Inference of Temporal Graph Neural Networks on FPGA. IEEE Transactions on Parallel and Distributed Systems 2024; 36(3):502–519.

274. Zhang B, Kannan R, Busart C, Prasanna V K. Model Architecture Codesign for High-Performance and Energy-Efficient SAR ATR on FPGA. IEEE Transactions on Geoscience and Remote Sensing 2025; 63:1–18.

275. Zhang C, Stepputtis S, Sycara K, Xie Y. Enhancing Vision-Language Few-Shot Adaptation with Negative Learning. 2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) 2025; 5905–5915.

276. Xu Z, Ye T, Kannan R, Prasanna V K. FAST: FPGA Acceleration of Fully Homomorphic Encryption with Efficient Bootstrapping. Proceedings of the 2025 ACM/SIGDA International Symposium on Field Programmable Gate Arrays 2025; 115–126.

277. Iftekhar A, Ruschel R, Kumar S, You S, Manjunath B S. DDS: Decoupled Dynamic Scene-Graph Generation Network. 2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) 2025; 9670–9680.

278. Wan Z et al. Sigma: Siamese Mamba Network for Multi-Modal Semantic Segmentation. 2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) 2025; 1734–1744.

279. Mao L, Warnell G, Stone P, Biswas J. pacer: Preference-Conditioned All-Terrain Costmap Generation. IEEE Robotics and Automation Letters 2025; 10(5):4572–4579.

280. Noseworthy M et al. FORGE: Force-Guided Exploration for Robust Contact-Rich Manipulation Under Uncertainty. IEEE Robotics and Automation Letters 2025; 10(5):4436–4443.

281. Zhao J et al. The Federation Strikes Back: A Survey of Federated Learning Privacy Attacks, Defenses, Applications, and Policy Landscape. ACM Computing Surveys 2025; 57(9):1–37.

282. Liu Y et al. AP Selection in Uplink Cell-Free Massive MIMO: An Unsupervised Heterogeneous GNN Approach. 2025 IEEE Wireless Communications and Networking Conference (WCNC) 2025; 1–6.

283. Paritosh P, Atanasov N, Martínez S. Distributed Bayesian Estimation in Sensor Networks: Consensus on Marginal Densities. IEEE Transactions on Network Science and Engineering 2025; 12(4):2848–2862.

284. Tang X, Panda A, Sehwag V, Mittal P. Differentially Private Image Classification by Learning Priors from Random Processes. Journal of Privacy and Confidentiality 2025; 15(1).

285. Song D, Liang J, Xiao X, Manocha D. VL-TGS: Trajectory Generation and Selection Using Vision Language Models in Mapless Outdoor Environments. IEEE Robotics and Automation Letters 2025; 10(6):5791–5798.

286. Mandal N, Khajenejad M, Martínez S. Distributed Task Allocation for Self-Interested Agents With Partially Unknown Rewards. IEEE Transactions on Automatic Control 2025; 70(9):6284–6291.

287. Meng Y, Iyer M A, Prasanna V K. An Acceleration Framework for Deep Reinforcement Learning Using Heterogeneous Systems. IEEE Transactions on Parallel and Distributed Systems 2025; 36(7):1401–1415.

288. Sivaprakasam M et al. SALON: Self-supervised Adaptive Learning for Off-road Navigation. 2025 IEEE International Conference on Robotics and Automation (ICRA) 2025; 16999–17006.

289. Seo M et al. PRESTO: Fast Motion Planning Using Diffusion Models Based on Key-Configuration Environment Representation. 2025 IEEE International Conference on Robotics and Automation (ICRA) 2025; 10861–10867.

290. Ho C et al. MapEx: Indoor Structure Exploration with Probabilistic Information Gain from Global Map Predictions. 2025 IEEE International Conference on Robotics and Automation (ICRA) 2025; 13074–13080.

291. Parikh D, Fein-Ashley J, Ye T, Kannan R, Prasanna V. ClusterViG: Efficient Globally Aware Vision GNNs via Image Partitioning. 2025 International Conference on Multimedia Computing, Networking and Applications (MCNA) 2025; 136–143.

292. Daiya D, Conover D, Bera A. COLLAGE: Collaborative Human-Agent Interaction Generation Using Hierarchical Latent Diffusion and Language Models. 2025 IEEE International Conference on Robotics and Automation (ICRA) 2025; 8203–8210.

293. Rivera C G et al. ConceptAgent: LLM-Driven Precondition Grounding and Tree Search for Robust Task Planning and Execution. 2025 IEEE International Conference on Robotics and Automation (ICRA) 2025; 8988–8995.

294. Gao P, Jose W J, Zhang H. Bandwidth-Adaptive Spatio-temporal Correspondence Identification for Collaborative Perception. 2025 IEEE International Conference on Robotics and Automation (ICRA) 2025; 15769–15775.

295. Yang T, Wang Y, Wei C, You S, Kuo C J. Efficient human object-interaction (EHOI) detection via interaction label coding and Conditional Decision. Computer Vision and Image Understanding 2025; 258:104390.

296. Shah K et al. AeroGen: Ground-to-Air Generalization for Action Recognition. 2025 IEEE 19th International Conference on Automatic Face and Gesture Recognition (FG) 2025; 1–10.

297. Abraham S J, Hauenstein J D, Scheirer W J. Towards Fair and Robust Face Parsing for Generative AI: A Multi-Objective Approach. 2025 IEEE 19th International Conference on Automatic Face and Gesture Recognition (FG) 2025; 1–10.

298. Lu C, Derakhshandeh K, Chaterji S. Improving Semi-Supervised Semantic Segmentation with Sliced-Wasserstein Feature Alignment and Uniformity. 2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) 2025; 20233–20243.

299. Reddy A et al. Video-ColBERT: Contextualized Late Interaction for Text-to-Video Retrieval. 2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) 2025; 19691–19701.

300. Zhou S et al. Feature4X: Bridging Any Monocular Video to 4D Agentic AI with Versatile Gaussian Feature Fields. 2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) 2025; 14179–14190.

301. Bresciani C et al. Preliminary Insights Into Resource-Constrained Neuro-Symbolic Causal Complex Event Processing. 2025 28th International Conference on Information Fusion (FUSION) 2025; 1–8.

302. Kowshik S A, Fisseler A, Srinivasa A R, Reddy J N. Physics informed Gaussian process models for real-time simulation of tire terrain interactions for off-road conditions. Journal of Terramechanics 2025; 120:101077.

303. Zhang C, Wan Z, Stepputtis S, Sycara K, Xie Y. Spectral-aware Global Fusion for RGB-Thermal Semantic Segmentation. 2025 IEEE International Conference on Image Processing (ICIP) 2025; 43–48.

304. Shah K, Sinha A, Reddy A, Roy A, Chellappa R. DIFFUSE2ADAPT: Controlled Diffusion for Synthetic-to-Real Domain Adaptation. 2025 IEEE International Conference on Image Processing (ICIP) 2025; 2235–2240.

305. Baber C, Cassenti D N. Comparing Static and Dynamic Decision Tasks Using Pupillometry: Expected Loss is More Important than Defined Reward. Proceedings of the Human Factors and Ergonomics Society Annual Meeting 2025.

Appendix B. Army Artificial Intelligence Innovation Institute (A2I2) Research Contributors

A systematic effort was undertaken to identify and disambiguate all authors and institutions connected to the A2I2 publications. This process produced a final set of 742 unique authors affiliated with 116 distinct institutions. Of these, 32 institutions were international, with the remainder (84) based in the United States. The institutions with the largest numbers of contributing authors included:

- Army Research Laboratory: 65 authors
- Purdue University: 55 authors
- Carnegie Mellon University: 50 authors
- University of Southern California: 50 authors
- University of California, Los Angeles: 45 authors
- The University of Texas at Austin: 35 authors
- Texas A&M University: 30 authors

Additional institutions are listed in this appendix. Several important considerations shaped the construction of this dataset:

- Author disambiguation: Multiple naming variations for the same individual were reconciled. For example, Celso de Melo, Celso M. de Melo, Celso Miguel de Melo, Celso M. De Melo, and Celso De Melo were determined to be the same person. This analysis was combined with comparing institutions on publications as well. For instance, there were multiple authors named Peng Gao with different institutions on papers; no visible connection between them was found—i.e., from moving between institutions (see cross-referencing below)—so they were determined to be separate authors with the same name.
- Author correction: In a few instances, the author last name was incorrectly spelled on the article. For example, Prof. Reza Langari's name was spelled "Lanagri" on a publication, Prof. Rama Chellappa's name was spelled "Chellapp" on another publication. When these mistakes happened, they were corrected.
- Affiliation corrections: In some cases, affiliations were incorrectly listed by publishers or authors. These were cross-checked and corrected. For example, ARL-affiliated authors sometimes used inconsistent identifiers such as DEVCOM Army Research Office, DEVCOM ARL West, or other variants. In a similar manner, many faculty and students may list their department or their laboratory in addition to university. These were standardized for consistency.
- Incomplete or incorrect attribution: In cases where initiatives were listed as af-

filiations,[*] or when institutions were provided without disambiguation across co-authors,[†] the affiliation is recorded as None. These are listed within academia as the most likely categorization. Fortunately, many of these names appeared on more than one article, so the institution from other articles were assumed to be correct one for the associated name.

- Cross-institutional mobility: Some co-authors were connected to funded institutions through prior employment or collaborative networks. Changes in affiliation were documented when possible, and ORCID IDs were used extensively to verify author identity and track institutional movement across time. In some cases, it was determined that the prior institution should be used rather than the affiliation listed on the paper because the timing indicates that the contribution likely occurred prior to the move.
- Cross-referencing sources: ORCID IDs (and links to personal webpages or Google Scholar), publisher metadata, and institutional rosters were cross-referenced to improve accuracy and resolve cases where the same name appeared across multiple institutions.

This process required extensive manual verification and investigative measures to ensure the accuracy of both author names and institutional affiliations. The resulting dataset captures not only the breadth of contributors but also highlights the collaborative networks that developed across academia, industry, government, and international institutions. The lists below were separated into multiple subsections: ARL, U.S. Academia, International Academia, Government, and FFRDCs (non-ARL), and Industry. For the last two categories, both domestic and international are listed together with international institutions having the country listed in parenthesis.

[*]This occurred once. An author listed "Advanced Distributed Learning (ADL) Initiative" without listing the institution.

[†]This occurred once. An arXiv article listed five authors and three institutions, but did not show which authors were associated with which institutions.

Army Research Laboratory

DEVCOM Army Research Laboratory (65): Mitchell Abrams, David M. Alexander, Robert St. Amant, Ahmed H. Anwar, Anthony L. Baker, Claire Bonial, Christoph C. Borel-Donohue, James Dean Brown, Carl Busart, Daniel N. Cassenti, Kevin Chan, Marshal Childers, Damon M. Conover, John A. D'Agostino, Venkateswara Dasari, James M. Dotterweich, Sungmin Eum, Ashley Foots, Billy Geerhart, Felix Gervits, Jason M. Gregory, Gregory M. Gremillion, James Z. Hare, André Harrison, Ethan Harrison, Cory J. Hayes, Ahmed Hemida, Cassidy Henry, Susan G. Hill, Shuowen Hu, Taylor Hudson, James Humann, Charles Kamhoua, Rajgopal Kannan, Lance Kaplan, Chad C. Kessens, Heesung Kwon, Stephanie M. Lukin, Celso Miguel de Melo, Jesse Milzman, Vinod K. Mishra, Mark Mittrick, Priya Narayanan, Carlos Nieto-Granda, Kimberly A. Pollard, Domenick Poster, Long P. Quang, Raghuveer Rao, Jean-Paul Reddinger, Brian Reily, Joe T. Rexwinkle, John T. E. Richardson, John G. Rogers, Brian M. Sadler, Aamodh Suresh, Ananthram Swami, Matthew Thielke, James Uplinger, Nikolas Vale, Clare R. Voss, Peng Wang, Garrett Warnell, Nicholas R. Waytowich, Noah Weston, Suya You

U.S. Academia

Arizona State University (2): Nicolò Michelusi, Zhe Xu
Auburn University at Montgomery (1): Tianran Chen
Brown University (1): Narek Harutyunyan
Caltech (1): Tinashe Handina
Carnegie Mellon University (50): Aditya Agarwal, Shubhra Aich, Sandip Aine, Joseph Campbell, Mateo Guaman Castro, Xianyi Cheng, Howie Choset, Manash Pratim Das, Siddharth Girdhar, Yupeng Han, Cherie Ho, Fahad Islam, Carlee Joe-Wong, Aaron M. Johnson, Seungchan Kim, Taejin Kim, Oliver Kroemer, Alex LaGrassa, Moon Young Lee, Steven Lee, Tabitha Edith Lee, Jacky Liang, Maxim Likhachev, Brady Moon, Shohin Mukherjee, Ramkumar Natarajan, Pradeep Ravikumar, Andrew Saba, Amrita S. Sawhney, Dhruv Mauria Saxena, Saumya Saxena, Sebastian Scherer, M. Yunus Seker, Mohit Sharma, Matthew Sivaprakasam, Simon Stepputtis, Katia Sycara, Samuel Triest, Shivam Vats, Anirudh Vemula, Zifu Wan, Sean J. Wang, Wenshan Wang, Steve Willits, Yaqi Xie, Hanlan Yang, Peng Yin, Silong Yong, Ce Zhang, Jialiang Zhao

Case Western Reserve University (2): Fei Shao, Xusheng Xiao

Cornell University (2): Angelique Taylor, Daniel Farkash

Duke University (2): Maria Gorlatova, Michael K. Reiter

Fordham University (1): Maryam Saleki

George Mason University (17): Dibyendu Das, Aniket Datar, Zechen Hu, Manshi Limbu, Yuanjie Lu, Mohammad Nazeri, Chenhui Pan, Amirreza Payandeh, Anuj Pokhrel, Amir Hossain Raj, Daigo Shishika, Daeun Song, Gregory J. Stein, Junzhe Wang, Xuan Wang, Xuesu Xiao, Yanlin Zhou

Georgia Institute of Technology (1): Judy Hoffman

Johns Hopkins University (19): Wele Gedara Chaminda Bandara, Rama Chellappa, Benjamin Van Durme, Deepti Hegde, Reno Kriz, Chun Pong Lau, Alexander Martin, Kenton Murray, Vishal M. Patel, Yasiru Ranasinghe, Aniket Roy, Kate Sanders, Lalithkumar Seenivasan, Anshul Shah, Ketul Shah, Mathias Unberath, Eugene Yang, Andrew Yates, Bryanna Yeh

Los Gatos High School (1): Sarthak Jain

Massachusetts Institute of Technology (8): Javier Alonso–Mora, Krishna Murthy Jatavallabhula, Michael Noseworthy, Hui Ren, Nicholas Roy, Seiji Shaw, Joshua B. Tenenbaum, Antonio Torralba

Michigan State University (1): Shaunak D. Bopardikar

New York University (1): Quanyan Zhu

Northeastern University (2): Shahin Shahrampour, Youbang Sun

Pennsylvania State University (4): Isaiah Adu, Daniel Cullina, Feiran Jia, Peng Liu

Portland State University (1): Atul Ingle

Princeton University (23): Sihui Dai, Kaixuan Huang, Sanjeev R. Kulkarni, Saeed Mahloujifar, Prateek Mittal, Roberto Morabito, Milad Nasr, Hung T. Nguyen, Ashwinee Panda, Xiangyu Qi, Vikash Sehwag, Virat Shejwalkar, Liwei Song, Xinyu Tang, Alexander Valtchanov, Sameer Wagh, Jiachen T. Wang, Mengdi Wang, Tianhao Wang, Tong Wu, Chong Xiang, Tinghao Xie, David Sommer

Purdue University (55): Mustafa Abdallah, Vaneet Aggarwal, Sheikh Shams Azam, Saurabh Bagchi, Ruqi Bai, Akhil Bandarupalli, Aniket Bera, Adithya Bhat, Christopher G. Brinton, Timothy N. Cason, Somali Chaterji, Wei Chen, Mung Chiang, Wonwoong Cho, Yun-Wei Chu, Ahaan Dabholkar, Divyanshu Daiya, Kasra Derakhshandeh, Mai Elkady, Bhargav Ganguly, Ziyu Gong, Dong-Jun Han, Seyyedali Hosseinalipour, David I. Inouye, Aniket Kate, Kwang Taik Kim, Murat

Kocaoğlu, Sean Kulinski, Jayoung Lee, Jiacheng Li, Ninghui Li, Xiang Li, Jim Lim, Xiaojun Lin, David J. Love, Chen-Yi Lu, Akash Melachuri, Preeti Mukherjee, Manish Nagaraj, David Nickel, Joseph Pappas, Rohit Parasnis, Qiang Qiu, Atul Sharma, Shreyas Sundaram, Dhruv Swarup, Satyavrat Wagle, Pengcheng Wang, Rui Wang, Su Wang, Xiaoqian Wang, Ran Xu, Joshua C. Zhao, Zeyu Zhou, Yihan Zou

San Diego State University (1): Tingting Tang

Stanford University (6): Leonidas Guibas, Peter Henderson, Jeong Joon Park, Despoina Paschalidou, Yijia Weng, Gordon Wetzstein

Texas A&M University (30): Madhu Areti, Anant Bhamri, Austin Carroll, Kenny Chour, Swaroop Darbha, Kameron Eves, Andrew Fisseler, Santosh Ganji, Swaminathan Gopalswamy, Zohaib Hasnain, Suhas A. Kowshik, Lucas W. Krakow, P. R. Kumar, Reza Langari, Jaekyung Jackie Lee, Tianqi Li, Tzu-Hsiang Lin, Tao Liu, Anthony Medellin, Rishi Phatak, Sivakumar Rathinam, Junuthula N. Reddy, Vamseedhar R. Reddyvari, Akshay Sarvesh, Srinivas Shakkottai, Dylan A. Shell, Arun R. Srinivasa, John Valasek, David van Wijk, Zhenyu Wu

The Ohio State University (18): Wei-Lun Chao, Shijie Chen, Xiang Deng, Xiang Deng, Yu Gu, Jihyung Kil, Lingbo Mo, Parinaz Naghizadeh, Vardaan Pahuja, Chan Hee Song, Samuel Stevens, Yu Su, Huan Sun, Boshi Wang, Clayton Washington, Jiaman Wu, Kai Zhang, Boyuan Zheng

The University of Texas at Austin (35): Joydeep Biswas, Rohan Chandra, Shenghui Chen, Zhiwen Fan, David Fridovich-Keil, Zichao Hu, Mustafa Karabag, Haresh Karnan, Bo Liu, Antonio M. López, Luisa Mao, Anirudh Nair, Aryaman Singh Samyal, Mikyung Seo, Peter Stone, Yoonchang Sung, Rishabh Saumil Thakkar, Adam J. Thorpe, Ufuk Topcu, Zhangyang Wang, Dejia Xu, Zifan Xu, Shuicheng Yan, Yue Yu, Yuke Zhu, Zehao Zhu, Pranav Atreya, Gauraang Dhamankar, Kang Hou, Fulin Jiang, Shuozhe Li, Sadegh Rabiee, Kavan Singh Sikand, Adam Uccello, Zizhao Wang

The University of Texas at Dallas (1): Murat Kantarcıoğlu

The University of Texas at El Paso (1): Deepak K. Tosh

Tulane University (1): Xin Hu

University at Buffalo (2): Weifeng Su, Chen Wang

University of California, Berkeley (14): Chih-Yuan Chiu, Peng Gao, Joseph E. Gonzalez, Jingqi Li, Xiaoyuan Liu, Michael W. Mahoney, Fernando Palafox, Kannan Ramchandran, Somayeh Sojoudi, Dawn Song, Linyue Song, Costas

J. Spanos, Claire J. Tomlin, Yaoqing Yang

University of California, Irvine (1): Congmiao Li

University of California, Los Angeles (45): Parth Agrawal, Shreeram Athreya, Yunhao Ba, Xinzhu Bei, Zachary Berger, Tejas Bharadwaj, Oyku Deniz Bozkurt, Julian de Gortari Briseno, Gianna Brown, Chethan Chinder Chandrappa, Haoran Chang, Pradyumna Chari, Xiangning Chen, Minhao Cheng, Gaofeng Dong, Blake Gella, Liying Han, Cho-Jui Hsieh, Bahram Jalali, Laleh Jalilian, Jeya Vikranth Jeyakumar, Sicheng Jiang, Achuta Kadambi, Narayanan Kannan, Tian Yu Liu, Varan Mehra, Parth Patwa, Arnold Pfahnl, Vwani Roychowdhury, Ankur Sarker, Stefano Soatto, Mani Srivastava, Akira Suzuki, Chinmay Talegaonkar, Niranjan Vaddi, Brian Wang, Ruochen Wang, Zhen Wang, Alex Wong, Ethan Yang, Yanchao Yang, Howard Zhang, Shuwang Zhang, Shijie Zhou, Yiming Zhou

University of California, San Diego (7): Nikolay Atanasov, Scott Brown, Mohammad Khajenejad, Nirabhra Mandal, Sónia Martínez, Parth Paritosh, Laurel D. Riek

University of California, Santa Barbara (9): Ivan Arevalo, ASM Iftekhar, Satish Kumar, Bangalore S. Manjunath, R. Austin McEver, Ekta Prashnani, Raphael Ruschel, Yu-Xiang Wang, Yuqing Zhu

University of California, Santa Cruz (1): Cihang Xie

University of Chicago (5): Arjun Nitin Bhagoji, Christian Cianfarani, Wenxin Ding, Ben Y. Zhao, Haitao Zheng

University of Colorado Boulder (5): Dirk Grunwald, Sangtae Ha, Jinsung Lee, Hyoyoung Lim, Sandesh Dhawaskar Sathyanarayana

University of Delaware (1): Xing Gao

University of Denver (1): Christopher Reardon

University of Illinois Urbana-Champaign (9): Girish Chowdhary, Mateus V. Gasparino, Vitor Akihiro Hisano Higuti, Bo Li, Yixiao Liu, Sayan Mitra, Arun N. Sivakumar, Huy Tran, Andrés Eduardo Baquero Velasquez

University of Maryland, College Park (9): Shuvra S. Bhattacharyya, Mohamed Elnoor, Peng Gao, Yaesop Lee, Jing Liang, Ming C. Lin, Dinesh Manocha, Adarsh Jagan Sathyamoorthy, Yi-Ting Shen

University of Massachusetts Amherst (6): Chuang Gan, Peng Gao, Amir Houmansadr, Ahmed Jaafar, Williard Joshua Jose, Hao Zhang

University of Michigan (3): Dmitry Berenson, Peter Mitrano, Elvin Yang

University of Montana (1): Anh Tuan Nguyen

University of Notre Dame (6): Sophia J. Abraham, Jonathan D. Hauenstein, Jeffery Kinnison, Kehelwala Dewage Gayan Maduranga, Zachary M. Miksis, Walter J. Scheirer

University of Pittsburgh (1): Micah Nye

University of Southern California (50): Ömer Faruk Akgül, Salman Avestimehr, Hong-Shuo Chen, Gangda Deng, Ahmed Roushdy Elkordy, Yahya H. Ezzeldin, Jacob Fein-Ashley, Hongyu Fu, Hongyu Fu, Luis Antonio García, Hamza Ghani, Neelesh Gupta, Rakshith Jayanth, C.-C. Jay Kuo, Shouyi Li, Joseph Lin, Ruiyuan Lin, Yi-Chien Lin, Yuhong Liu, Yuhuai Liu, Yuan Meng, Dhruv Parikh, Viktor K. Prasanna, Cauligi S. Raghavendra, Mozhdeh Rouhsedaghat, Ishika Singh, Bingjie Tang, Jesse Thomason, Ye Tian, Ta-Yang Wang, Xinyu Wang, Xu Wang, Yifan Wang, Yun-Cheng Wang, Chengwei Wei, Sachini Wickramasinghe, Samuel Wiggins, Sasindu Wijeratne, Zhihan Xu, Tsung-Shan Yang, Yijing Yang, Tian Ye, Hanqing Zeng, Bingyi Zhang, Kaitai Zhang, Pengmiao Zhang, Youning Zhang, Zhiwei Zhang, Hongkuan Zhou, Zhiruo Zhou

University of Toronto (2): Qiao Gu, Florian Shkurti

University of Waterloo (1): Wenhu Chen

University of Wisconsin–Madison (16): Matthew Dutson, Michael Gleicher, Mohit Gupta, Mark Han, Sacha Jungerman, Yin Li, Yiquan Li, Yingyu Liang, Sicheng Mo, Fangzhou Mu, Zhenmei Shi, Carter Sifferman, Junyi Wei, Zhuoyan Xu, Felipe Gutierrez, Andreas Velten

Vanderbilt University (2): Garrison L. H. Johnston, Nabil Simaan

Virginia Tech (9): Tyler Cody, Daniel Doyle, Christopher T. Franck, Laura Freeman, Ruoxi Jia, Erin Lanus, Jeremy Ogorzalek, Sami Wood, Yi Zeng

International Academia

Birla Institute of Technology and Science, Pilani (1, India): Aditya Parandekar

Dalian University of Technology (2, China): Yuhao Wang, Pingping Zhang

Delft University of Technology (1, Netherlands): Lasse Peters

École Polytechnique Fédérale de Lausanne (6, Switzerland): Maksym Andriushchenko, Edoardo Debenedetti, Nicolas Flammarion, Andrei Ardelean, Claudio Bruschini, Edoardo Charbon

ETH Zurich (4, Switzerland): David Dao, Nezihe Merve Gürel, Frances Ann Hubis, Ce Zhang

Imperial College London (3, United Kingdom): Leo Ardon, Daniel Furelos-Blanco, Alessandra Russo
Indian Institute of Science (2, India): Sarat Chandra Bobbili, Parimal Parag
Indian Institute of Technology Kharagpur (1, India): Parv Maheshwari
Indian Institute of Technology Madras (1, India): Abhiram Rao Gorle
Korea Advanced Institute of Science and Technology (4, Korea): Yoonyoung Cho, Beomjoon Kim, Do-Yeon Kim, Jaekyun Moon
Kyung Hee University (1, Korea): Minseok Choi
Nanjing University (2, China): Zheng Qin, Fengyuan Xu
National University of Singapore (1, Singapore): Jeric Lew
Sapienza University of Rome (1, Italy): Aymeric Barbin
Seoul National University (3, Korea): Junseon Kim, Jongyun Lee, Kyunghan Lee
Southeast University (1, China): Zhengming Zhang
Technion – Israel Institute of Technology (2, Israel): Eyal Kushilevitz, Oren Salzman
Ulsan National Institute of Science and Technology (1, Korea): Youngbin Im
Universität Innsbruck (1, Austria): Daniel Woods
Université de Montréal (5, Canada): Kirsty Ellis, Ali Kuwajerwala, Sacha Morin, Liam Paull, Bipasha Sen
University College London (1, United Kingdom): Simon Julier
University of Birmingham (1, United Kingdom): Chris Baber
University of Brescia (9, Italy): Pietro Baroni, Christian Bresciani, Federico Cerutti, Marco Cominelli, Alfonso Gerevini, Massimiliano Giacomin, Francesco Gringoli, Luca Lavazza, Lorenzo Nodari
University of Dschang (4, Cameroon): Clémentin Tayou Djamegni, Fabrice Mvah, Achile Leonel Nguemkam, Vianney Kengne Tchendji
University of Technology, Sydney (1, Australia): Graeme Best
University of Tübingen (2, Germany): Francesco Croce, Matthias Hein
Vrije Universiteit (1, Netherlands): Lucia Donatelli
Zhejiang University (3, China): Ping He, Yiming Li, Boxin Wang

Government and FFRDCs (non-ARL)

Air Force Research Laboratory (4): Trevor Bihl, Erik Blasch, Nandini Iyer, Felix J. Knutson
DARPA (1): Matthew Marge
Defense Threat Reduction Agency (2): Kenneth Evensen, Frank Skirlo
Department of Energy (1): G. Jeremy Leong
DEVCOM Ground Vehicle Systems Center (2): David Grabowsky, Dariusz Mikulski
JHU Applied Physics Laboratory (15): Andrew R. Badger, Meghan Booker, Erin C. Butler, Grayson Byrd, Martha Cervantes, Tyler Feldman, David A. Handelman, Emma Holmes, Kapil D. Katyal, Bethany Kemp, Rohita Mocharla, William Paul, Arun Reddy, Corban G. Rivera, Aurora Schmidt
National Science Foundation (1): Erion Plaku
NATO CMRE (2, Italy): Paolo Braca, Leonardo M. Millefiori
Qatar Computing Research Institute (1, Qatar): Issa Khalil
USC Institute for Creative Technologies (4): Ron Artstein, Anton Leuski, David Traum, Setareh Nasihati Gilani

Industry

Adobe Research (6): Shubham Agarwal, Kanak Mahadik, Subrata Mitra, Anup Rao, Shiv Kumar Saini, Md Mehrab Tanjim
Algorand Foundation (2): Fabrice Benhamouda, Tal Rabin
Amazon (2): Dimitrios Dimitriadis, Arash Nourian
Amazon Alexa AI (1, United Kingdom): Murat Şensoy
Aurora Innovation (1): J. Andrew Bagnell
Booz Allen Hamilton (1): Francisca Alba-Sánchez
Cisco Research (2): Hugo Latapie, Jayanth Srinivasa
Cisco Systems (1): Shikhar Suryavansh
DiDi AI Labs (1): Xiaocheng Tang
Everyday Robots (1): Sven Mikael Persson
Fibertek (1): Joshua D. Gray
Google (17): Boqing Gong, Mingxing Tan, Li Zhang, Krzysztof Choromański, Anthony Francis, Roy Frostig, Dmitry Kalashnikov, Edward Lee, Carolina Parada, Vikas Sindhwani, Sumeet Singh, Jie Tan, Stephen Tu, Jake Varley, Fei Xia, Peng

Xu, Tingnan Zhang

Google DeepMind (1, United Kingdom): Christopher A. Choquette-Choo

Hoku Labs (1): Leila Takayama

IBM Research (1): Pin-Yu Chen

Intel Corporation (1): Mahesh A. Iyer

Meta (3): Christopher Leung, Jianbo Li, Yinglong Xia

Microsoft Research (1): Shruti Tople

NVIDIA (13): Iretiayo Akinola, Valts Blukis, Dieter Fox, Animesh Garg, Ankit Goyal, Ankur Handa, Arsalan Mousavian, Yashraj Narang, Fábio Ramos, Holger R. Roth, Jonathan Tremblay, Bowen Wen, Danfei Xu

Oxbridge Health (1): Arushi Sinha

Soar Technology (2): Jeffrey D. Craighead, Ben Purman

Vanguard Group (1): Dhagash Mehta

Zillow Group (1): Naji Khosravan

Appendix C. Army Artificial Intelligence Innovation Institute (A2I2) Paper Citations

Appendix Table C-1 lists the citation histories of individual papers as recorded in OpenAlex. Each entry is organized by Digital Object Identifier (DOI), which also provides a direct hyperlink to the corresponding publication.

The table reports the number of citations received by each paper for each year from 2020 through 2025. A running total (Sum) is provided across all years, along with the percentile placement of the paper within its global field reference set (%). In addition, two binary indicators identify whether the paper falls within the top 1% or top 10% of cited papers worldwide for its field and publication year. In addition to raw citation counts, the table includes a citation-normalized percentile from OpenAlex. This measure situates each paper within the distribution of works of the same type, year, and research field, and in many cases also the same venue class (e.g., journal or conference proceedings). For example, a paper in the 99th percentile has received more citations than 99% of comparable works published in the same context. Unlike averages, which can be distorted by a small number of extremely highly cited papers, the percentile emphasizes relative standing within the relevant distribution. This makes it a useful complement to raw citation counts, particularly in fields with skewed, fat-tailed citation patterns. Note that OpenAlex does not compute this percentile for arXiv entries.

Because of the large number of entries, the table is formatted to span multiple pages in landscape orientation for readability. This presentation allows for a straightforward year-by-year view of citation growth while also enabling comparison across papers based on cumulative counts and percentile placement.

Table C-1. **Papers and citation histories (OpenAlex).**

No.	DOI	2020	2021	2022	2023	2024	2025	Sum	%	Top-10%	Top-1%
1	10.1109/icra48891.2023.10161317	0	0	3	48	195	196	442	100.0%	YES	YES
2	10.1109/iccv51070.2023.00280	0	0	0	15	118	152	285	99.9%	YES	YES
3	10.2478/popets-2021-0011	3	25	34	42	40	53	197	99.5%	YES	YES
4	10.14778/3342263.3342637	16	21	19	27	32	29	147	97.1%	YES	
5	10.1109/icra57147.2024.10610243	0	0	0	0	40	87	127	100.0%	YES	YES
6	10.48550/arxiv.2010.09670	2	47	11	16	28	11	115			
7	10.1109/icde51399.2021.00024	0	7	17	30	22	26	102	99.0%	YES	YES
8	10.1109/cvpr52733.2024.02048	0	0	0	0	19	73	92	100.0%	YES	YES
9	10.1109/icme51207.2021.9428361	0	11	28	25	16	10	90	98.2%	YES	
10	10.48550/arxiv.2003.10595	3	10	22	27	14	8	84			
11	10.1109/iccv48922.2021.01251	0	1	16	27	27	8	79	96.4%	YES	
12	10.1609/aaai.v38i19.30150	0	0	0	1	26	41	68	99.3%	YES	YES
13	10.1109/cvpr46437.2021.01635	0	4	20	15	16	8	63	97.8%	YES	
14	10.1109/lra.2021.3090023	1	4	12	10	16	17	60	95.4%	YES	
15	10.1145/3460120.3484757	0	1	7	13	15	19	55	96.1%	YES	
16	10.1109/lra.2022.3193464	0	0	0	13	25	16	54	99.4%	YES	YES
17	10.1109/wacv56688.2023.00338	0	0	0	5	20	24	49	94.3%	YES	
18	10.1109/icra48891.2023.10160856	0	0	0	3	22	22	47	97.4%	YES	
19	10.1109/tnsm.2022.3216326	0	0	3	12	13	18	46	99.1%	YES	YES
20	10.48550/arxiv.2103.12051	0	19	5	7	6	6	44			
21	10.56553/popets-2022-0072	0	0	4	6	15	18	43	96.9%	YES	
22	10.48550/arxiv.2310.03693	0	0	0	3	19	16	38			

Continued on next page

Table C-1. Papers and citation histories (OpenAlex) (continued).

#	DOI	2020	2021	2022	2023	2024	2025	Sum	%	Top-10%	Top-1%
23	10.1109/wacv61041.2025.00176	0	0	0	0	0	37	37	100.0%	YES	YES
24	10.48550/arxiv.2005.10884	0	2	11	11	11	1	36			
24	10.1016/j.robot.2022.104132	0	3	5	5	14	9	36	94.2%	YES	
26	10.1109/icra46639.2022.9812316	0	1	5	8	16	5	35	91.3%	YES	
27	10.1109/icra46639.2022.9811648	0	0	1	7	17	9	34	99.3%	YES	YES
28	10.1109/cvpr52688.2022.01993	0	0	3	14	11	5	33	99.2%	YES	YES
28	10.48550/arxiv.2206.10341	0	1	2	15	11	4	33			
30	10.1109/tnnls.2022.3190364	0	0	0	7	9	14	30	95.0%	YES	
30	10.1038/s42256-023-00662-0	0	0	1	0	9	20	30	97.1%	YES	
32	10.1109/cvpr46437.2021.00096	0	1	6	12	10	0	29	92.5%	YES	
32	10.1109/iros47612.2022.9982060	0	0	3	5	13	8	29	98.9%	YES	YES
34	10.1109/icra48891.2023.10160583	0	0	0	4	10	14	28	95.3%	YES	
34	10.48550/arxiv.2306.10012	0	0	0	1	20	7	28			
34	10.1145/3560830.3563730	0	0	1	5	12	10	28	94.6%	YES	
37	10.1109/sp54263.2024.00030	0	0	0	0	6	21	27	99.0%	YES	YES
37	10.1109/jsac.2024.3365901	0	0	0	0	10	17	27	98.9%	YES	YES
37	10.48550/arxiv.2306.06070	0	0	0	2	14	11	27			
40	10.1109/cvpr52729.2023.00033	0	0	0	2	14	9	25	98.0%	YES	
40	10.1109/sp46215.2023.10179319	0	0	0	6	4	15	25	96.4%	YES	
42	10.1145/3580804	0	0	0	3	13	8	24	94.2%	YES	
42	10.1109/spw54247.2022.9833884	0	0	2	9	10	3	24	93.3%	YES	
44	10.1109/mcom.005.2200925	0	0	0	2	14	7	23	95.9%	YES	

Continued on next page

Table C-1. Papers and citation histories (OpenAlex) (continued).

#	DOI	2020	2021	2022	2023	2024	2025	Sum	%	Top-10%	Top-1%
44	10.1109/icstw55395.2022.00031	0	0	4	3	7	9	23	97.2%	YES	
46	10.1145/3492321.3519577	0	0	0	9	7	6	22	90.2%	YES	
46	10.48550/arxiv.2110.08324	0	0	3	8	7	4	22			
46	10.1109/satml54575.2023.00024	0	0	0	3	13	6	22	95.7%	YES	
46	10.1109/cvpr52729.2023.01689	0	0	0	2	7	13	22	94.6%	YES	
46	10.1109/tmc.2023.3274708	0	0	0	3	9	10	22	93.0%	YES	
51	10.1109/icra48891.2023.10161268	0	0	0	2	12	7	21	90.6%	YES	
51	10.1109/icra57147.2024.10610079	0	0	0	0	7	14	21	98.0%	YES	
51	10.1109/icra57147.2024.10611265	0	0	0	0	5	16	21	96.9%	YES	
54	10.1109/lra.2024.3511409	0	0	0	0	1	19	20	97.8%	YES	
55	10.1145/3466772.3467030	0	1	5	4	8	1	19	90.8%	YES	
55	10.48550/arxiv.2108.09135	0	0	0	8	11	0	19			
55	10.1109/icra48891.2023.10160416	0	0	0	2	7	10	19	92.1%	YES	
58	10.1007/978-3-031-72658-3_19	0	0	0	0	2	14	16	98.3%	YES	YES
58	10.1109/iros47612.2022.9981259	0	0	2	1	8	5	16	84.2%		
58	10.1109/icra46639.2022.9811828	0	0	2	1	7	6	16	83.5%		
61	10.1109/cvpr52733.2024.01014	0	0	0	0	3	12	15	99.0%	YES	YES
61	10.1109/sp46214.2022.9833591	0	0	0	5	8	2	15	90.6%	YES	
61	10.1561/116.00000126	0	0	2	3	7	3	15	89.5%	YES	
61	10.1145/3583780.3614769	0	0	0	2	7	6	15	92.8%	YES	
65	10.1109/tmc.2023.3331690	0	0	0	0	7	7	14	92.7%	YES	
65	10.1109/icra48891.2023.10161366	0	0	0	2	7	5	14	91.7%	YES	

Continued on next page

Table C-1. Papers and citation histories (OpenAlex) (continued).

#	DOI	2020	2021	2022	2023	2024	2025	Sum	%	Top-10%	Top-1%
65	10.1109/tnet.2023.3262482	0	0	0	2	8	4	14	92.3%	YES	
65	10.1109/icra46639.2022.9811575	0	0	2	5	7	0	14	81.7%		
69	10.1145/3565287.3610278	0	0	0	1	9	3	13	92.0%	YES	
69	10.1109/icra57147.2024.10611710	0	0	0	0	1	12	13	95.7%	YES	
69	10.48550/arxiv.2112.06274	0	0	0	9	2	2	13			
69	10.1109/cvpr52688.2022.00256	0	0	1	3	3	6	13	81.4%		
73	10.1109/ssrr56537.2022.10018758	0	0	0	2	6	4	12	82.3%		
73	10.1109/cvpr52729.2023.01297	0	0	0	2	8	2	12	87.0%		
73	10.1016/j.robot.2023.104442	0	0	0	3	7	2	12	95.6%	YES	
73	10.1109/icra48506.2021.9561881	0	0	2	1	5	4	12	74.9%		
73	10.1145/3469116.3470010	0	2	3	4	0	3	12	76.3%		
73	10.1109/lra.2024.3457369	0	0	0	0	5	7	12	96.5%	YES	
79	10.1109/tip.2023.3245994	0	0	0	3	4	4	11	85.4%		
79	10.1109/lsp.2021.3103130	0	2	3	2	2	2	11	83.8%		
79	10.1109/cvpr52729.2023.00387	0	0	0	3	7	1	11	90.4%	YES	
79	10.1016/j.cose.2023.103696	0	0	0	0	3	8	11	95.8%	YES	
79	10.1145/3579856.3582836	0	0	0	2	5	4	11	90.2%	YES	
84	10.1109/iccv51070.2023.01551	0	0	0	2	2	6	10	85.0%		
84	10.1109/cvpr52733.2024.00507	0	0	0	0	1	9	10	97.2%	YES	
84	10.1109/cvpr52733.2024.02667	0	0	0	0	2	8	10	94.0%	YES	
84	10.1145/3724113	0	0	0	0	0	10	10	99.7%	YES	YES
84	10.1109/access.2023.3299235	0	0	0	2	3	5	10	84.3%		

Continued on next page

Table C-1. Papers and citation histories (OpenAlex) (continued).

#	DOI	2020	2021	2022	2023	2024	2025	Sum	%	Top-10%	Top-1%
84	10.1109/tpami.2021.3071289	2	3	0	0	1	1	10	70.7%		
90	10.1609/socs.v15i1.21761	0	0	0	4	5	0	9	77.6%		
90	10.1109/lra.2022.3157544	0	0	1	3	5	0	9	75.8%		
90	10.23919/acc55779.2023.10155841	0	0	0	2	4	3	9	86.2%		
90	10.56553/popets-2022-0112	0	0	2	6	1	0	9	83.4%		
90	10.1145/3469259.3470487	0	1	0	2	3	3	9	88.1%		
95	10.1117/12.2663936	0	0	0	1	6	1	8	80.5%		
95	10.1109/cvpr52688.2022.01475	0	0	0	4	2	2	8	75.5%		
95	10.1109/hpec58863.2023.10363455	0	0	0	0	4	4	8	93.5%	YES	
95	10.1109/radar54928.2023.10371090	0	0	0	0	3	5	8	87.4%		
95	10.1109/mra.2023.3322920	0	0	0	0	2	6	8	85.3%		
95	10.1109/iros58592.2024.10801650	0	0	0	0	0	8	8	92.2%	YES	
95	10.1109/icra57147.2024.10611319	0	0	0	0	3	5	8	91.8%	YES	
95	10.1109/mra.2024.3427891	0	0	0	0	1	7	8	91.9%	YES	
103	10.1109/icra57147.2024.10610790	0	0	0	0	2	5	7	90.2%	YES	
103	10.1109/icra57147.2024.10610668	0	0	0	0	1	6	7	84.0%		
103	10.1109/ccta48906.2021.9658603	0	0	2	3	1	1	7	83.1%		
103	10.1109/icra48891.2023.10161297	0	0	0	0	4	3	7	82.6%		
103	10.1109/icip51287.2024.10647347	0	0	0	0	1	6	7	90.5%	YES	
103	10.1109/vcip53242.2021.9675333	0	1	2	3	1	0	7	75.4%		
109	10.1109/icassp49357.2023.10096564	0	0	0	2	3	1	6	82.7%		
109	10.1109/icra48891.2023.10161006	0	0	0	1	3	2	6	75.7%		

Continued on next page

Table C-1. Papers and citation histories (OpenAlex) (continued).

#	DOI	2020	2021	2022	2023	2024	2025	Sum	%	Top-10%	Top-1%
109	10.48550/arxiv.2107.06929	0	2	0	1	3	0	6	73.4%		
109	10.1109/ipdps54959.2023.00062	0	0	0	0	6	0	6	83.1%		
109	10.48550/arxiv.2305.01639	0	0	0	1	3	2	6			
109	10.48550/arxiv.2104.09425	0	0	1	1	3	1	6			
109	10.1109/lra.2025.3559822	0	0	0	0	0	6	6	99.9%	YES	YES
109	10.1109/cvpr52733.2024.01791	0	0	0	0	0	6	6	88.7%	YES	
109	10.1016/j.patrec.2021.05.009	0	3	2	0	1	0	6	67.7%		
109	10.1109/iccv48922.2021.01021	0	0	2	1	3	0	6	94.6%	YES	
109	10.1109/bigdata59044.2023.10386982	0	0	0	0	3	3	6	81.7%		
120	10.23919/apsipaasc55919.2022.9980233	0	0	1	3	1	0	5	65.2%		
120	10.1145/3658644.3670326	0	0	0	0	1	4	5	90.7%	YES	
120	10.1609/aaai.v35i7.16765	0	0	0	2	2	1	5	61.8%		
120	10.1609/icaps.v33i1.27249	0	0	0	0	4	1	5	72.3%		
120	10.1109/lra.2025.3551637	0	0	0	0	0	5	5	98.2%	YES	YES
120	10.1109/icra48506.2021.9562058	0	0	3	1	1	0	5	70.6%		
120	10.1109/ipdps57955.2024.00087	0	0	0	0	2	3	5	89.8%	YES	
120	10.1109/sbac-pad59825.2023.00012	0	0	1	0	2	2	5	75.2%		
120	10.1561/116.00000008	0	0	0	2	3	0	5	56.2%		
120	10.1109/lra.2022.3174366	0	0	1	1	1	2	5	65.5%		
120	10.48550/arxiv.2206.09868	0	0	0	2	3	0	5			
120	10.48550/arxiv.2006.01952	3	2	0	0	0	0	5			
120	10.1109/hpec58863.2023.10363615	0	0	0	0	0	5	5	81.9%		

Continued on next page

Table C-1. Papers and citation histories (OpenAlex) (continued).

#	DOI	2020	2021	2022	2023	2024	2025	Sum	%	Top-10%	Top-1%
120	10.48550/arxiv.2205.13613	0	0	1	3	1	0	5			
120	10.1145/3564289	0	0	0	0	2	3	5	66.4%		
120	10.1109/globecom48099.2022.10000962	0	0	0	1	4	0	5	66.0%		
120	10.15607/rss.2020.xvi.025	0	3	0	0	1	1	5	55.7%		
137	10.1109/case49997.2022.9926704	0	0	0	0	2	2	4	60.3%		
137	10.48550/arxiv.2308.12439	0	0	1	0	1	2	4			
137	10.1109/tiv.2024.3363177	0	0	1	1	2	0	4	76.2%		
137	10.1109/cvpr52733.2024.01164	0	0	0	0	1	3	4	88.1%	YES	
137	10.1109/lra.2024.3433207	0	0	0	0	2	2	4	93.0%	YES	
137	10.1109/hpec62836.2024.10938499	0	0	0	0	0	4	4	88.8%	YES	
137	10.1109/lra.2024.3520919	0	0	0	0	0	4	4	80.8%		
137	10.1109/icra57147.2024.10611488	0	0	0	0	2	2	4	82.5%		
137	10.1109/icra57147.2024.10611475	0	0	0	0	0	4	4	82.7%		
137	10.1109/jstsp.2024.3374591	0	0	0	0	1	3	4	85.4%	YES	
137	10.1145/3590140.3592848	0	0	0	0	3	1	4	77.2%		
137	10.1109/icra55743.2025.11128862	0	0	0	0	0	4	4	99.7%	YES	YES
137	10.1109/mmsp55362.2022.9949325	0	0	1	2	1	0	4	60.2%		
137	10.1109/lra.2023.3290533	0	0	0	1	3	0	4	82.3%		
137	10.48550/arxiv.2308.15709	0	0	0	1	2	1	4			
137	10.1109/dcoss-iot58021.2023.00044	0	0	0	0	4	0	4	88.2%		
137	10.23919/acc55779.2023.10156477	0	0	0	1	1	2	4	73.5%		
137	10.1109/whispers61460.2023.10431091	0	0	0	1	2	1	4	79.4%		

Continued on next page

Table C-1. Papers and citation histories (OpenAlex) (continued).

#	DOI	2020	2021	2022	2023	2024	2025	Sum	%	Top-10%	Top-1%
137	10.1109/srds55811.2022.00013	0	0	0	0	2	2	4	70.2%		
156	10.1109/hpec58863.2023.10363610	0	0	0	0	1	2	3	65.7%		
156	10.1109/hpec58863.2023.10363595	0	0	0	0	0	3	3	69.1%		
156	10.1109/tnet.2021.3092008	1	0	0	1	1	0	3	59.2%		
156	10.1109/hpec58863.2023.10363567	0	0	0	0	3	0	3	75.4%		
156	10.2514/6.2023-0128	0	0	0	1	0	2	3	80.0%		
156	10.1109/iscas48785.2022.9937539	0	0	1	0	1	1	3	53.7%		
156	10.1109/radarconf2458775.2024.10549364	0	0	0	0	1	2	3	90.5%	YES	
156	10.1145/3649153.3649187	0	0	0	0	0	3	3	76.0%		
156	10.1145/3649153.3649191	0	0	0	0	1	2	3	83.8%		
156	10.2514/6.2024-0230	0	0	0	0	1	2	3	75.3%		
156	10.48550/arxiv.2403.00871	0	0	0	0	1	2	3			
156	10.1016/j.autcon.2024.105393	0	0	0	0	1	2	3	90.5%	YES	
156	10.1145/3706628.3708879	0	0	0	0	0	3	3	98.2%	YES	
156	10.1109/cvpr52733.2024.01790	0	0	0	0	2	1	3	84.9%		
156	10.1109/infocom53939.2023.10229047	0	0	0	0	2	1	3	71.1%		
156	10.48550/arxiv.2205.13616	0	0	0	1	0	2	3			
156	10.1109/iros47612.2022.9981459	0	0	0	1	2	0	3	53.4%		
156	10.48550/arxiv.2209.10780	0	0	0	3	0	0	3			
156	10.1109/cvpr52729.2023.02107	0	0	0	0	1	2	3	63.5%		
156	10.1109/icra48891.2023.10160382	0	0	0	0	1	2	3	68.5%		
156	10.1117/12.2618686	0	0	0	2	0	1	3	67.0%		

Continued on next page

Table C-1. Papers and citation histories (OpenAlex) (continued).

#	DOI	2020	2021	2022	2023	2024	2025	Sum	%	Top-10%	Top-1%
156	10.1609/socs.v15i1.21753	0	0	0	1	2	0	3	58.3%		
156	10.1109/icra57147.2024.10610619	0	0	0	0	1	2	3	82.9%	YES	
156	10.1109/cdc49753.2023.10383423	0	0	0	1	0	2	3	78.4%		
156	10.1109/ipdps57955.2024.00082	0	0	0	0	0	3	3			
181	10.48550/arxiv.2104.02297	0	1	1	0	0	0	2			
181	10.1109/msec.2024.3407593	0	0	0	1	0	1	2	80.3%	YES	
181	10.23919/fusion59988.2024.10706320	0	0	0	0	0	2	2	79.5%	YES	
181	10.1109/hpec62836.2024.10938519	0	0	0	0	0	2	2	82.0%		
181	10.1145/3649153.3649212	0	0	0	0	0	2	2	68.3%		
181	10.48550/arxiv.2212.09736	0	0	0	1	1	0	2			
181	10.48550/arxiv.2210.10275	0	0	0	0	1	1	2			
181	10.1109/access.2023.3260405	0	0	0	1	0	1	2	52.6%		
181	10.1109/mcna65829.2025.11124368	0	0	0	0	0	2	2	92.4%	YES	
181	10.48550/arxiv.2307.04942	0	0	0	0	1	1	2			
181	10.48550/arxiv.2206.05573	0	0	0	2	0	0	2	64.6%		
181	10.48550/arxiv.2012.01693	0	1	0	1	0	0	2			
181	10.48550/arxiv.2011.04627	0	0	0	2	0	0	2			
181	10.1016/j.eswa.2024.125906	0	0	0	0	0	2	2	81.4%		
181	10.48550/arxiv.2411.12844	0	0	0	0	2	0	2	81.0%		
181	10.1109/wacv61041.2025.00576	0	0	0	0	0	2	2	99.1%	YES	YES
181	10.1109/fccm60383.2024.00017	0	0	0	0	1	1	2	71.6%		
181	10.1109/lcsys.2024.3410149	0	0	0	0	1	1	2	80.6%		

Continued on next page

Table C-1. Papers and citation histories (OpenAlex) (continued).

#	DOI	2020	2021	2022	2023	2024	2025	Sum	%	Top-10%	Top-1%
181	10.1016/j.jvcir.2024.104058	0	0	0	0	1	1	2	75.9%		
181	10.1109/icnc59896.2024.10556188	0	0	0	0	0	2	2	75.2%		
181	10.1109/fccm60383.2024.00018	0	0	0	0	0	2	2	81.8%	YES	
181	10.1109/icra57147.2024.10610129	0	0	0	0	1	1	2	75.2%		
181	10.1007/978-3-031-56255-6_3	0	0	0	0	2	0	2	84.3%	YES	
181	10.1109/radarconf2458775.2024.10548693	0	0	0	0	1	1	2	78.0%		
181	10.1109/irc59093.2023.00055	0	0	0	1	0	1	2	70.7%		
181	10.48550/arxiv.2309.15302	0	0	0	0	1	1	2	82.6%		
181	10.1145/3610661.3617166	0	0	0	0	0	2	2	70.3%		
181	10.1109/globecom46510.2021.9685893	0	0	0	1	1	0	2	53.4%		
181	10.1109/cdc49753.2023.10384097	0	0	0	0	0	2	2	70.3%		
181	10.1609/socs.v16i1.27295	0	0	0	2	0	0	2	66.2%		
181	10.1117/12.2662753	0	0	0	1	1	0	2	65.8%		
181	10.1609/socs.v16i1.27297	0	0	0	1	1	0	2	66.2%		
181	10.1109/ur57808.2023.10202423	0	0	0	0	2	0	2	55.5%		
214	10.1109/cvpr52734.2025.01834	0	0	0	0	0	1	1	90.3%	YES	
214	10.48550/arxiv.2210.03780	0	0	0	1	0	0	1			
214	10.1109/iros45743.2020.9341475	0	1	0	0	0	0	1	46.5%		
214	10.1016/j.cviu.2025.104390	0	0	0	0	0	1	1	87.7%	YES	
214	10.1109/icra55743.2025.11128268	0	0	0	0	0	1	1	95.6%	YES	
214	10.1109/tpds.2025.3566766	0	0	0	0	0	1	1	94.0%	YES	
214	10.29012/jpc.910	0	0	0	0	0	1	1	91.8%	YES	

Continued on next page

Table C-1. Papers and citation histories (OpenAlex) (continued).

#	DOI	2020	2021	2022	2023	2024	2025	Sum	%	Top-10%	Top-1%
214	10.1109/tpds.2024.3521897	0	0	0	0	0	1	1	74.1%		
214	10.1109/wcnc61545.2025.10978823	0	0	0	0	0	1	1	79.9%	YES	
214	10.1109/lra.2025.3549645	0	0	0	0	0	1	1	82.5%	YES	
214	10.1109/icra48506.2021.9561439	0	0	0	0	1	0	1	51.8%		
214	10.48550/arxiv.2403.12210	0	0	0	0	1	0	1			
214	10.1145/3624062.3624169	0	0	0	0	0	1	1	59.6%		
214	10.1109/access.2024.3513461	0	0	0	0	0	1	1	68.6%		
214	10.1007/978-3-031-47754-6_3	0	0	0	0	1	0	1	82.3%		
214	10.1109/techdefense59795.2023.10380847	0	0	0	0	0	1	1	57.0%		
214	10.48550/arxiv.2311.00815	0	0	0	0	1	0	1			
214	10.48550/arxiv.2402.15017	0	0	0	0	1	0	1			
214	10.1109/iccv51070.2023.00987	0	0	0	0	1	0	1	80.0%		
214	10.48550/arxiv.2104.08382	0	0	1	0	0	0	1			
214	10.1109/iros58592.2024.10802451	0	0	0	0	0	1	1	84.4%		
214	10.48550/arxiv.2207.02286	0	0	0	0	1	0	1			
214	10.48550/arxiv.2204.13779	0	0	0	0	0	1	1			
214	10.1109/ai4i51902.2021.00016	0	1	0	0	0	0	1	56.3%		
214	10.1117/12.2663960	0	0	0	0	0	1	1	61.0%		
214	10.1109/ro-man57019.2023.10309403	0	0	0	0	1	0	1	59.6%		
214	10.48550/arxiv.2206.12728	0	0	0	1	0	0	1			
214	10.1117/12.2663827	0	0	0	0	1	0	1	45.6%		
214	10.1109/cvpr52733.2024.01369	0	0	0	0	0	1	1	60.4%		

Continued on next page

Table C-1. Papers and citation histories (OpenAlex) (continued).

#	DOI	2020	2021	2022	2023	2024	2025	Sum	%	Top-10%	Top-1%
214	10.1177/15485129241271178	0	0	0	0	1	0	1	70.1%		
214	10.1117/12.3013530	0	0	0	0	0	1	1	55.8%		
214	10.1117/12.3013906	0	0	0	0	0	1	1	80.2%		
214	10.1109/tpds.2024.3466891	0	0	0	0	0	1	1	60.6%		

List of Symbols, Abbreviations, and Acronyms

2D	two-dimensional
3D	three-dimensional
3DGS	3D Gaussian Splatting
A2I2	Army AI Innovation Institute
AAAI	Association for the Advancement of Artificial Intelligence
ACM	Association for Computing Machinery
AI	artificial intelligence
APPL	Adaptive Planner Parameter Learning
ARO	Army Research Office
ARL	Army Research Laboratory
BAA	Broad Agency Announcement
C2	command and control
C4I	command, control, communications, computers, and intelligence
CDC	Conference on Decision and Control
CNN	convolutional neural network
CRA	Collaborative Research Alliance
CPU	central processing unit
CVPR	Conference on Computer Vision and Pattern Recognition
DEVCOM	U.S. Army Combat Capabilities Development Command
DOI	Digital Object Identifier
DOW	Department of War

DTIC	Defense Technical Information Center
EO/IR	electro-optical/infrared
EW	electronic warfare
FFRDC	Federally Funded Research and Development Center
GPU	graphics processing unit
GPS	global positioning system
HPEC	High Performance Extreme Computing Conference
ICCV	International Conference on Computer Vision
ICIP	International Conference on Image Processing
IEEE	Institute of Electrical and Electronics Engineers
ICRA	International Conference on Robotics and Automation
ID	identifier
IPDPS	International Parallel and Distributed Processing Symposium
IROS	International Conference on Intelligent Robots and Systems
ISR	intelligence, surveillance, and reconnaissance
JHU	Johns Hopkins University
LiDAR	light detection and ranging
LLM	large language model
LPTM	large pretrained model
MDO	Multi-Domain Operations
ML	machine learning
NLP	natural language processing
ORCID	Open Researcher and Contributor ID

OSCTI	Open-Source Cyber Threat Intelligence
PI	principal investigator
RF	radio frequency
RGB-D	red-green-blue-depth
SPIE	International Society for Optics and Photonics
SSD	self-supervised outlier detection
SSL	successive subspace learning
SWaP	size, weight, and power
SWaPT	size, weight, and power, time
TBQL	Threat Behavior Query Language
UAV	unmanned aerial vehicle
UGV	unmanned ground vehicle
USC	University of Southern California
WACV	Winter Conference on Applications of Computer Vision

1 (PDF)	DEFENSE TECHNICAL INFORMATION CTR DTIC OCA
1 (PDF)	DEVCOM ARL FCDD RLB CI TECH LIB