

Autonomous Resilient Mesh Infrastructure for Distributed Edge Intelligence (ARMI-DEI)

NSF Program Solicitation: CISE / Robust Intelligence (RI) – Core Small

Lead Institution: Vertex Institute of Technology

Principal Investigator: Dr. Elena Rostova (Vertex)

Co-PI: Dr. Marcus Vance (Meridian University)

Proposed Duration: 36 Months (Oct 2026 – Sep 2029)

Total Requested Budget: \$598,420

Target Track: RI Small Collaborative Research

1 Project Overview and Research Motivation

Modern autonomous cyber-physical systems—ranging from autonomous aerial swarm networks operating in communication-denied environments to real-time smart grid power balancing systems—increasingly depend on decentralized decision-making at the network edge. Traditional centralized cloud architecture suffers from brittle network dependencies, excessive latency (> 100 ms), and catastrophic single-point failures when operating in dynamic, adversarial, or resource-constrained settings. While distributed machine learning (e.g., Federated Learning) has advanced edge analytics, current paradigms assume continuous topological connectivity and stable link bandwidth. They lack formal safety guarantees when network links undergo sudden degradation or Byzantine failures.

This project introduces **ARMI-DEI** (Autonomous Resilient Mesh Infrastructure for Distributed Edge Intelligence), a novel framework that unifies *neural-symbolic reasoning*, *dynamic topological consensus*, and *formal verification primitives* directly on heterogeneous edge hardware. By embedding lightweight formal logic engines into deep neural architectures, ARMI-DEI enables autonomous agents to perform provably correct, localized inference while dynamically reorganizing inter-node communication meshes under severe network partitioning.




Figure 1: ARMI-DEI system architecture illustrating autonomous edge mesh reconfiguration during cloud disconnection.

1.1 Key Research Objectives

The overarching goal of ARMI-DEI is to transform vulnerable edge deployments into self-healing, safety-verified intelligent swarms. The research is structured around three core scientific objectives:

- Objective 1: Differentiable Neural-Symbolic Reasoning Layers (Thrust 1):** Formulate mathematical primitives that inject Temporal Logic (TL) specifications into gradient-based neural network training, ensuring zero-violation safety bounds during local edge inference.
- Objective 2: Topological Graph Consensus under Intermittent Connectivity (Thrust 2):** Develop decentralized, asynchronous consensus algorithms capable of maintaining state consistency across dynamic mesh topologies with link loss up to 60%.
- Objective 3: Hardware-Aware Micro-Verification and Field Testbed (Thrust 3):** Realize sub-millisecond execution of formal safety constraints on embedded SoC architectures (Lumina Jetson, RISC-V edge TPU) and validate performance across a 15-node autonomous ground rover fleet.

2 Results from Prior NSF Support

PI: Dr. Elena Rostova – NSF Award IIS-2104892: “Autonomous Edge Perception in Dynamic Environments” (\$499,820; Period: 09/01/2021 – 08/31/2024).

Intellectual Merit: Under this prior award, the PI’s lab investigated lightweight graph neural networks (GNNs) for multi-robot perception. Key outcomes included: (1) A spatio-temporal GNN architecture that reduced perception inference latency by 42% on low-power Jetson Xavier platforms; (2) An adaptive feature compression algorithm yielding a $10\times$ reduction in inter-agent wireless bandwidth requirements without loss of target classification accuracy; and (3) 8 peer-reviewed publications in IEEE Transactions on Robotics (T-RO), IEEE ICRA, and NeurIPS.

Broader Impacts: The project supported the dissertation research of 3 PhD students (2 from underrepresented groups in STEM) and 6 undergraduate researchers. The PI developed a new graduate course, *ECE 6840: Edge Artificial Intelligence*, which reached over 140 students across three academic years. Software artifacts, including the *EdgeSense-100K* multi-sensor dataset, were made publicly available on GitHub (> 450 stars, > 1,200 downloads), fostering open science in the robotics research community.

Relationship to Proposed Work: While IIS-2104892 focused on *perception* and feature compression, the proposed ARMI-DEI project addresses the orthogonal and critical challenges of *reasoning*, *formal safety verification*, and *decentralized consensus* under severe network degradation.

3 Research Plan and Technical Methodology

The research plan is organized into three tightly integrated technical thrusts spanning 36 months.

3.1 Thrust 1: Differentiable Neural-Symbolic Reasoning Layers

Deep neural networks excel at perceptual pattern recognition but lack explainability and deterministic guarantees. Conversely, symbolic logic provides rigorous safety specifications but suffers from combinatorial state explosion. Thrust 1 bridges this gap by creating differentiable loss formulations from Signal Temporal Logic (STL) specifications.

Let $s(t) \in \mathbb{R}^d$ denote the state trajectory of an autonomous edge node at time t . An STL property φ defines safety constraints (e.g., "always maintain distance $d_{\min} > 1.5$ m from obstacles"). We define a continuous, smooth quantitative robustness measure $\rho(\varphi, s, t)$ such that $\rho(\varphi, s, t) > 0 \iff s \models \varphi$.

$$\mathcal{L}_{\text{total}}(\theta) = \mathcal{L}_{\text{task}}(f_{\theta}(x), y) + \lambda \cdot \text{SoftPlus}\left(-\rho(\varphi, f_{\theta}(x), t)\right) \quad (1)$$

By backpropagating through Equation 1, the model parameters θ are optimized directly toward satisfaction of symbolic safety properties without requiring expensive SMT solver calls at runtime.

3.2 Thrust 2: Topological Graph Consensus under Intermittent Connectivity

Edge networks in contested or emergency response scenarios experience frequent topology changes and link drops. We model the communication topology as a time-varying directed graph $\mathcal{G}(t) = (\mathcal{V}, \mathcal{E}(t))$. We formulate an asynchronous matrix completion method for edge state synchronization:

Phase	Algorithmic Primitive	Key Innovation	Latency Target
Phase I	Local State Representation	Compact Bit-Vector Encoding (≤ 128 bytes)	< 0.5 ms
Phase II	Asynchronous Gossip Consensus	Bounded-Error Matrix Completion	< 3.2 ms
Phase III	Safety Verification Check	Embedded SMT Micro-Checker	< 1.0 ms
Phase IV	Control Actuation	Fallback Safe Policy Execution	< 0.2 ms

Table 1: Operational breakdown of the ARMI-DEI edge decision pipeline per control loop cycle.

3.3 Thrust 3: Embedded Micro-Verification and Testbed Validation

To evaluate ARMI-DEI under realistic physical dynamics, we will deploy the software pipeline on a 15-node physical testbed comprising custom ground rovers equipped with Lumina Jetson Orin Nano modules and RISC-V edge co-processors. Testbed experiments will take place at the Synapse Field Station facility, simulating urban search-and-rescue and micro-grid energy balancing.

💡 INTELLECTUAL MERIT HIGHLIGHT

Intellectual Merit Summary: ARMI-DEI fundamentally advances the state-of-the-art in robust intelligence by establishing the mathematical foundations for differentiable neural-symbolic reasoning. It provides the first unified framework capable of guaranteeing formal safety bounds (STL satisfaction) while maintaining asynchronous multi-agent consensus over lossy, dynamic edge topologies.

4 Project Schedule, Management, and Milestones

The project will be co-led by PI Dr. Elena Rostova (Vertex Institute; expert in neural-symbolic systems and edge AI) and Co-PI Dr. Marcus Vance (Meridian University; expert in formal verification and distributed control). Bi-weekly virtual synchronization meetings and semi-annual joint workshops will ensure seamless collaboration between both institutions.

Milestone	Description / Deliverable	Year 1	Year 2	Year 3	Lead
M1.1	Mathematical formulation of differentiable STL loss	■			Vertex
M1.2	Open-source PyTorch neural-symbolic toolkit release		■		Vertex
M2.1	Asynchronous gossip consensus protocol implementation	■			Meridian
M2.2	Network fault injection benchmarks (> 50% link loss)		■		Meridian
M3.1	Deployment of 15-node Jetson Orin rover testbed		■		Joint
M3.2	Comprehensive field demonstration & benchmark data release			■	Joint

Table 2: Project milestone timeline across the 36-month grant period.

5 Broader Impacts

The broader impacts of ARMI-DEI span STEM education, broadening participation, technology transfer, and critical infrastructure resilience.

5.1 Integration with STEM Education and Curriculum Development

The PIs will create an inter-institutional project-based course, *Autonomous Edge Systems*, co-taught at Vertex Institute and Meridian University. The course will feature hands-on labs utilizing the open-source ARMI-DEI software framework, enabling students to program miniature autonomous vehicles in virtual and physical environments.

5.2 Broadening Participation of Underrepresented Groups

Partnering with the *Meridian STEM Inclusion Network*, the project will recruit at least 4 undergraduate researchers annually from underrepresented backgrounds through a funded 10-week summer REU (Research Experiences for Undergraduates) program. PIs will actively mentor candidates, providing structured pathways to graduate education.

5.3 Open Science, Software Release, and Industry Transfer

All code, model weights, hardware schematics, and benchmarking datasets produced by ARMI-DEI will be released under permissive open-source licenses (MIT/Apache 2.0) on GitHub. PIs will engage with industry advisory boards at Nexa Systems and Apex Autonomous Infrastructure to translate research findings into commercial edge computing standards.

BROADER IMPACTS HIGHLIGHT

Broader Impacts Summary: ARMI-DEI impacts broader society by fortifying critical infrastructure (smart grids, emergency response swarms) against communication failures and cyber-disruptions. It broadens STEM participation through targeted REU internships, open-access educational modules, and robust public software release.

6 Summary of Intellectual Merit and Broader Impacts

Intellectual Merit: ARMI-DEI establishes a rigorous mathematical foundation uniting neural-symbolic reasoning, formal Temporal Logic verification, and asynchronous mesh consensus. It resolves fundamental trade-offs between deep learning adaptability and formal safety guarantees in dynamic, communication-restricted environments.

Broader Impacts: The proposed work delivers tangible societal benefits by safeguarding mission-critical edge deployments, training a diverse next-generation STEM workforce, providing open-source tools to the academic community, and transferring technology to industrial partners.