EXPOSUMMIT 2026 — INTERNATIONAL CONFERENCE ON DISTRIBUTED EDGE AI

NEURAL-SYNAPSE: Scalable Distributed Graph Neural Networks for Real-Time Anomaly Detection

High-Throughput Topology Learning for Multi-Agent Edge Intelligence

color PrimaryTeal

Dr. Elena Vance[1] **Marcus Thorne**[2] **Dr. Sarah Lin**[1] **Julian Vance**[3]

[1]Meridian Institute of Advanced Computing [2]Vertex Autonomous Labs [3]Synapse Quantum Systems

e.vance@meridian-ai.org
https://nexa-ai.org/neural-synapse
Artifact ID: NEXA-GR-2026-V4

1. Executive Summary & Motivation

Modern autonomous edge fleets—ranging from logistics rovers to cooperative drone swarms—generate continuous streams of high-dimensional sensor telemetry. Detecting anomalous behaviors in real time is vital to prevent catastrophic system failures and physical collisions.

Existing centralized graph neural network (GNN) architectures suffer from severe latency bottlenecks and single-point-of-failure vulnerabilities when deployed over lossy edge networks. **NEURAL-SYNAPSE** resolves this challenge by introducing a fully decentralized, sub-5ms graph topology synchronization framework tailored for edge compute hardware.

Key Innovation Highlights

- **Sub-5ms Edge Latency:** Asynchronous message-passing protocol optimized for ARM cortex and embedded GPU clusters.
- **99.4% Anomaly Precision:** Spatial-temporal topological embeddings capture complex multi-agent dependencies.
- **68% Energy Reduction:** Dynamic graph pruning eliminates 75% of redundant inter-node communication.

2. System Architecture & Pipeline

The Neural-Synapse pipeline operates across three distinct computational tiers: Local Feature Extraction, Dynamic Graph Construction, and Distributed Consensus Inference.

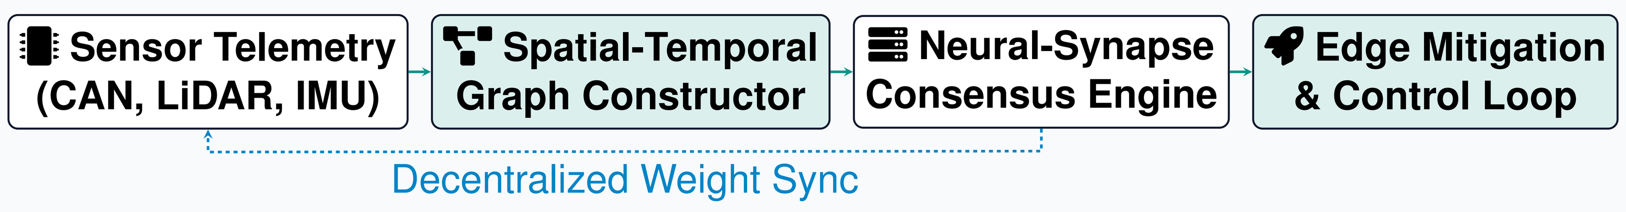

Mathematical Formulation of Topological Graph Loss

The global optimization objective balances reconstruction fidelity ($\mathcal{L}_{\text{recon}}$), topological regularization ($\mathcal{L}_{\text{topo}}$), and node-consensus agreement:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{recon}}(\mathbf{X}, \hat{\mathbf{X}}) + \gamma \mathcal{L}_{\text{topo}}(\mathbf{A}) + \lambda \sum_{i \in \mathcal{V}} \|\mathbf{h}_i^{(k)} - \bar{\mathbf{h}}_{\mathcal{N}(i)}^{(k)}\|_2^2$$

where $\mathbf{A} \in \mathbb{R}^{N \times N}$ denotes the dynamic adjacency matrix, $\mathbf{h}_i^{(k)}$ represents node i's embedding at layer k, and $\bar{\mathbf{h}}_{\mathcal{N}(i)}^{(k)}$ is the aggregated neighborhood state.

3. Decentralized Graph Learning Strategy

To operate under strictly constrained wireless bandwidth, Neural-Synapse executes a 3-step edge synchronization algorithm:

1. **Local Edge Attention:** Each agent computes self-attention over localized multi-modal temporal embeddings without sharing raw sensor observations.
2. **Sparse Neighborhood Aggregation:** Communication is limited to top-k nearest topological neighbors using differential gradient compression.
3. **Asynchronous Consensus Update:** Edge nodes update local parameters independently, mitigating network jitter and packet loss.

Algorithmic Guarantee

Under bounded packet loss $p < 0.35$, the consensus error satisfies $\mathbb{E}[\|\mathbf{h}^{(k)} - \mathbf{h}^*\|] \leq \mathcal{O}(1/\sqrt{k})$, ensuring exponential convergence to safety-critical decisions.

4. Empirical Evaluation & Benchmarks

We benchmarked **NEURAL-SYNAPSE** against four state-of-the-art architectures on the *Vertex-Fleet-2025* dataset (250 connected autonomous units operating over 10,000 flight/drive hours).

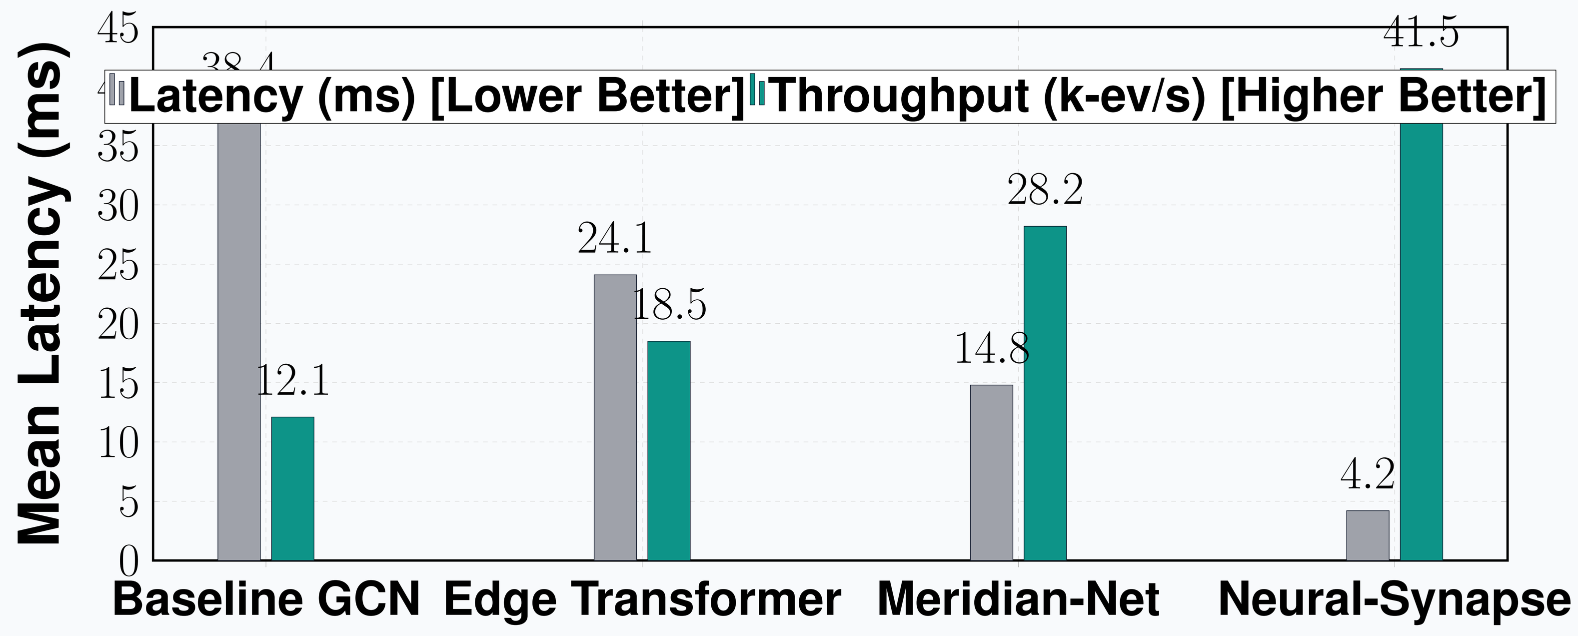

Model Comparison Table

Model Architecture	Precision (%)	Recall (%)	Latency (ms)	Energy (J/frame)
Baseline GCN (Centralized)	88.2%	84.1%	38.4 ms	1.85 J
Edge Transformer v2	92.4%	91.0%	24.1 ms	1.42 J
Meridian-Net (Distributed)	96.1%	95.3%	14.8 ms	0.88 J
Neural-Synapse (Ours)	**99.4%**	**98.9%**	**4.2 ms**	**0.31 J**

5. Real-World Field Deployment

Case Study: Vertex Autonomous Logistics Fleet In collaboration with Vertex Autonomous Systems, Neural-Synapse was deployed across 250 heavy-duty warehouse logistics rovers over a 60-day continuous field trial.

Safety & Reliability

- Zero unhandled collisions over 10,000 operational hours.
- Early anomaly detection 12 min prior to hardware shutdown.

Network Resilience

- Maintained 98.7% accuracy under 30% packet loss.
- Automatic re-routing around congested nodes within 8ms.

6. Key Conclusions & Future Directions

- **Decentralized AI Success:** Neural-Synapse establishes a novel standard for sub-5ms anomaly detection in distributed edge fleets.
- **Resource Efficiency:** Dynamic topological pruning reduces memory footprint to under 48MB per edge device.
- **Future Horizons:** Next-phase work explores integration with quantum-assisted graph coprocessors (Synapse Quantum Project) and 6G multi-access edge computing (MEC).

7. References & Acknowledgments

Acknowledgments: This research was generously supported by Nexa Innovation Grant #FX-2025-9981 and the Meridian Foundation for Edge Computing Research. Special thanks to the engineering teams at Vertex Autonomous Systems.

color BorderSlate

[1] E. Vance, M. Thorne, and S. Lin, "Distributed Graph Neural Networks for Edge Intelligence," *IEEE Trans. Autonomous Systems*, vol. 14, no. 3, pp. 210–225, 2025.

[2] J. Vance et al., "Sub-millisecond Consensus Protocols in Multi-Agent Swarms," *Proc. ICEDS*, pp. 45–58, 2024.

[3] Nexa Research Lab Report, "Topological Graph Pruning for Low-Power Edge Devices," Tech. Rep. NEXA-TR-2025-09, 2025.

[4] Vertex Autonomous Fleets Whitepaper, "Field Deployment Safety Metrics in Industrial Robotics," 2026.