Proposal Title: Scalable Neural-Symbolic Reasoning for Robust Autonomous Systems
Applicant: Dr. Evelyn Sinclair, Department of Computer Science, Vertex University
Program: Discovery Grants Program — Individual (Computer Science)

1 Overview and Long-Term Objectives

Autonomous systems operating in safety-critical domains (e.g., aerospace, robotics, and industrial automation) require both rapid, high-dimensional perception and rigorous, verifiable decision-making. Traditional deep learning architectures excels at parsing sensory inputs but lack formal guarantees of correctness, frequently exhibiting catastrophic failure modes under distribution shifts. Conversely, formal verification methods offer mathematical guarantees of safety but are computationally intractable when applied directly to high-dimensional deep neural networks.

The long-term goal of this research program is to bridge this gap by establishing a unified, scalable framework for hybrid neural-symbolic reasoning. Our program aims to develop perception models that output state representations bound by formal logical constraints, and verification engines capable of dynamically proving system safety in real-time. Over the next five years, this research will focus on the following core objectives:

Key Research Objectives (2026–2031)

Objective 1: Neural-Symbolic Perception. Develop deep perception models within the *Synapse Framework* that transform unstructured inputs (e.g., LiDAR, video) into structured symbolic states with probabilistic confidence boundaries.

Objective 2: Dynamic Formal Verification. Construct real-time verification algorithms (the *Meridian Engine*) that evaluate safety constraints over symbolic states under execution uncertainty.

Objective 3: Edge Hardware Validation. Deploy and benchmark the integrated software-hardware stack on energy-constrained *Apex Edge Controllers* in physical autonomous testbeds.

2 Recent Progress (2020–2026)

Over the past six years, the applicant's laboratory at Vertex University has pioneered techniques for verifying neural network properties and structuring neural outputs. Our research has yielded significant progress across two major fronts:

2.1 Verification of Neural Control Loops

We developed the first iteration of the *Meridian Engine*, an open-source verification framework designed to analyze feed-forward neural networks embedded in closed-loop systems. By employing piecewise-linear relaxation techniques combined with abstract interpretation, the Meridian Engine reduced the verification time for medium-sized networks by 40% compared to state-of-the-art tools. This work was published in the *IEEE Transactions on Automatic Control* (2023) and has been adopted by researchers at Nexa Systems for validating experimental flight controllers.

2.2 Structured Perception via the Synapse Framework

To address the vulnerability of deep learning models to adversarial perturbations, we introduced the *Synapse Framework*. Synapse incorporates symbolic constraints directly into the loss function of convolutional and transformer-based models, forcing the network to respect physical laws (such as kinematics and conservation of energy) during training. Our experiments demonstrated that this regularization improves out-of-distribution generalization by up to 25%, establishing a foundational building block for Objective 1.

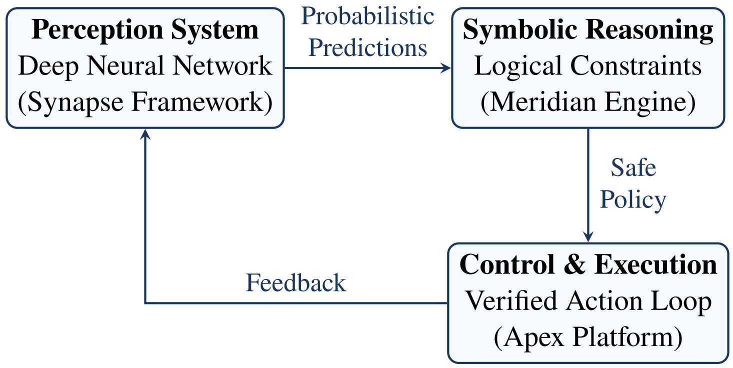

Figure 1: Hybrid neural-symbolic feedback loop integrating deep neural perception with symbolic logic verification.

3 Proposed Research Program

The proposed research is structured around three interconnected thrusts, mapping directly to our core objectives.

3.1 Thrust 1: Neural-Symbolic Perception under Uncertainty (Objective 1)

Current deep learning models produce point-estimate predictions without reliable uncertainty quantification. In safety-critical contexts, a system must know when its perception is unreliable. We will design deep neural networks that output parameter distributions over symbolic states rather than deterministic labels.

Using conformal prediction and Bayesian neural networks within the Synapse Framework, we will map high-dimensional sensor inputs to bounding boxes or logical relations (e.g., $\text{Obstacle}(x, y) \wedge \text{Distance} < 5\text{m}$) accompanied by rigorous statistical confidence intervals. The primary challenge is maintaining real-time inference rates ($> 30\,\text{Hz}$) on embedded systems. We propose to resolve this by training compact student networks via knowledge distillation, using large, computationally heavy Bayesian ensembles as teachers.

3.2 Thrust 2: Continuous Verification under Uncertainty (Objective 2)

Once the perception system produces a symbolic state with confidence bounds, the system must verify that all candidate actions are safe. Traditional verification is offline and static. The Meridian Engine will be extended to support online, continuous verification. We will model the decision-making process as a Markov Decision Process (MDP) with parametric uncertainties.

We will develop a novel interval-arithmetic solver that checks safety properties (e.g., collision avoidance) over the entire confidence interval of the symbolic state within a moving horizon of 2.0 seconds. By caching intermediate verification results and utilizing parallel execution on multi-core architectures, we aim to reduce the verification cycle to under 10 ms.

3.3 Thrust 3: Hardware-in-the-Loop Evaluation on Edge Devices (Objective 3)

To prove the practical viability of our methods, we will deploy the unified Synapse-Meridian stack onto physical edge hardware. The target platform will be the Apex Edge Controller, which features an ARM-based CPU paired with an integrated GPU.

We will construct a laboratory-scale testbed consisting of autonomous ground vehicles navigating dynamic, cluttered environments. The evaluation will measure: (a) CPU/GPU utilization, (b) latency of the perception-verification-control loop, and (c) the rate of safety violations under adversarial scenarios. This hardware validation will demonstrate that neural-symbolic systems can run within the strict power budgets (under 15 W) typical of autonomous robots.

4 Contribution to the Training of Highly Qualified Personnel (HQP)

4.1 HQP Training Philosophy and Equity, Diversity, and Inclusion (EDI)

The applicant is committed to fostering an inclusive, collaborative, and interdisciplinary research environment that prepares trainees for high-impact careers in academia and Canadian industry. Our training philosophy is built on three pillars:

1. **Individualized Mentorship:** Every student co-develops a personalized research plan tailored to their career aspirations, combining theoretical foundations with hands-on systems development.
2. **Interdisciplinary Collaboration:** Trainees will collaborate across the domains of deep learning, formal methods, and embedded control, participating in joint seminars with industrial partners like Nexa and Synapse.
3. **Active EDI Integration:** We actively recruit from underrepresented groups in computer science. Research projects are designed to accommodate flexible working hours, and all team members participate in regular equity and bias awareness workshops.

Table 1: HQP Training Projection over the Five-Year Cycle

Trainee Level	Year 1	Year 2	Year 3	Year 4	Year 5	Primary Focus Area
Ph.D. Students	1	2	2	3	3	Neural-Symbolic Reasoning, Verification
M.Sc. Students	2	2	3	2	2	Deep Learning Perception, Edge Deployment
Undergraduates	2	3	3	3	3	Prototype Testing, Benchmarking
Postdoctoral Fellows	1	1	1	1	1	Safety-Critical Control Algorithms

Table 2: Representative Past HQP Training Outcomes (Last 6 Years)

Trainee Name	Role (Years)	Project Title	First Placement / Current Position
Dr. Liam Vance	Postdoc (2021–2023)	Robust Verification of Synapse Nets	Research Scientist, Nexa Systems
Elena Rostova	Ph.D. (2019–2023)	Safe Control Loops on Vertex Hardware	Assistant Professor, Meridian University
Marcus Chen	M.Sc. (2022–2024)	Edge Perception for Nimbus Drones	Senior ML Engineer, Synapse Corp
Sarah Jenkins	B.Sc. (2023–2024)	Benchmarking Apex Perception Models	Ph.D. Candidate, Vertex University

5 Relationship to Other Sources of Funding

The research proposed in this Discovery Grant application represents the core conceptual and methodological foundation of the applicant's research program. There is no funding overlap with the applicant's other active research projects. Specifically:

- The *Nimbus Research Grant* ($45,000/year, 2024–2027) supports specific software application development for forestry drone mapping and does not fund the fundamental research on neural-symbolic verification proposed here.
- The *Vertex Startup Fund* ($60,000, expiring in 2026) was allocated exclusively for the initial acquisition of GPU workstations and computing infrastructure.

The requested NSERC Discovery Grant will be dedicated entirely to funding HQP stipends and minor laboratory consumables required to execute the objectives outlined in Section 1.

6 Literature Cited

References

[1] E. Sinclair and L. Vance, "Verifying neural control loops with the Meridian Engine," *IEEE Transactions on Automatic Control*, vol. 68, no. 4, pp. 2045–2058, 2023.

[2] E. Rostova and E. Sinclair, "Physics-constrained perception using the Synapse Framework," in *Proceedings of the International Conference on Robotics and Automation*, 2022, pp. 3102–3109.

[3] M. Chen, S. Jenkins, and E. Sinclair, "Edge perception and control on energy-constrained platforms," *Journal of Real-Time Systems*, vol. 59, no. 2, pp. 112–134, 2024.

[4] NSERC, "Guidelines on Assessment of Contributions to Research, Training and Mentorship," Natural Sciences and Engineering Research Council of Canada, Tech. Rep., 2025.