

ACADEMIC DEFENSE DOSSIER

VIVA VOCE PREPARATION GUIDE

Thesis Summary, Anticipated Questions & Defense Strategy

Thesis Title:

*Epistemic Calibration and Safety-Critical Risk Control
in Autonomous Drone Swarms*

Submitted for the degree of Doctor of Philosophy in Robotics & AI

Candidate: Alexander Mercer, M.Sc.
Department: Department of Autonomous Systems, Vertex Institute
Supervisor: Dr. Marcus Vance
Date of Viva: October 14, 2026
Venue: Senator Boardroom, Apex Science Centre

COMMITTEE PANEL

Prof. Sarah Jenkins	Vertex Institute of Technology (Internal Examiner)
Prof. Elena Rostova	Meridian Research Laboratory (External Examiner)
Dr. Aline Dubois	Synapse Labs (Committee Chair)

Contents

1	Introduction & Overview	1
1.1	The Structure of the Defense	1
1.2	Milestones and Timeline	1
2	Thesis Summary by Chapter	2
2.1	Chapter 1: Introduction & Epistemic Vulnerabilities	2
2.2	Chapter 2: Related Work & Traditional Controls	2
2.3	Chapter 3: Spatio-Temporal Navigation via SynapseNet	2
2.4	Chapter 4: Epistemic Calibration & The Apex Filter	2
2.5	Chapter 5: Hardware-in-the-Loop Validation	3
3	Anticipated Questions & Answers	4
4	Defence Strategy & Weak-points Table	6
4.1	Summary of Self-Criticism	6
	References	8

Introduction & Overview

The transition from a candidate to a doctoral degree holder is marked by the defense of the thesis, a process colloquially known as the *viva voce*. This document serves as a self-contained defense dossier compiled by the candidate to organize and outline key contributions, anticipated questions, defense tactics, and potential vulnerabilities.

1.1 The Structure of the Defense

The defense is structured around four primary milestones, designed to validate both the breadth of the candidate's domain knowledge and the depth of their original research contributions:

1. **Candidate Presentation:** A 20-minute concise overview of the thesis objectives, methodology, key findings, and practical implications.
2. **General Examination:** Broad questions regarding the field of study, establishing the candidate's mastery of the underlying literature.
3. **Chapter-by-Chapter Defense:** Detailed technical interrogation focusing on specific mathematical derivations, experimental setups, and interpretations.
4. **Contribution Review:** Critical discussion of the novelty, limitations, and future research paths emerging from the thesis work.

1.2 Milestones and Timeline

The preparation timeline is structured to allow progressive refinement of the defense arguments. Figure 1.1 visualizes the milestone timeline leading up to the defense date.



Figure 1.1. Viva Voce Preparation Timeline

By systematically evaluating the research questions, this dossier builds confidence and establishes a structured framework for addressing critical feedback during the viva.

Thesis Summary by Chapter

This chapter details the core contents of the thesis, organized by chapter. Each summary emphasizes the central problem, the proposed solution, and the corresponding empirical or theoretical findings.

2.1 Chapter 1: Introduction & Epistemic Vulnerabilities

Core Objective: Identify and formalize the risk of out-of-distribution (OOD) environments in deep reinforcement learning (DRL) controllers for unmanned aerial vehicles.

- **Key Problem:** Neural network controllers perform optimally under training distributions but exhibit catastrophic failures when sensor feeds degrade.
- **Major Contribution:** Formalized a risk framework based on epistemic uncertainty, mapping state space areas to probability distributions of control loss.

2.2 Chapter 2: Related Work & Traditional Controls

Core Objective: Critical review of Model Predictive Control (MPC) and robust control theory in comparison to deep learning alternatives.

- **Key Problem:** Classic MPC architectures lack the real-time processing capabilities required for high-velocity quadcopter operations in cluttered fields.
- **Major Contribution:** Proved that while MPC offers safety guarantees, the computational latency scales exponentially with obstacle count, making it unsuitable for swarm deployments.

2.3 Chapter 3: Spatio-Temporal Navigation via SynapseNet

Core Objective: Design and validation of a neural network model to handle trajectory generation in GPS-denied environments.

- **Key Problem:** Existing spatial networks ignore flight path histories, leading to drift in the presence of dynamic obstacles.
- **Major Contribution:** Developed **SynapseNet**, incorporating temporal convolutions with a spatial attention mask, which reduced flight path error by 22.4% under sensor degradation [1].

2.4 Chapter 4: Epistemic Calibration & The Apex Filter

Core Objective: Development of a real-time safety filter that guarantees safe trajectories under bounded uncertainty.

- **Key Problem:** DRL models do not provide deterministic safety guarantees.

- **Major Contribution:** Introduced the **Apex** Safety Filter (ASF), which uses control-barrier functions parameterized by real-time epistemic variance. This guarantees collision-free navigation given bounded sensor noise [2].

2.5 Chapter 5: Hardware-in-the-Loop Validation

Core Objective: Verify the proposed algorithms on hardware in simulated and physical test beds.

- **Key Problem:** Sim-to-real transfer errors often invalidate simulation results.
- **Major Contribution:** Validated the models in the **Nimbus** flight simulator and physical swarm testing at the Apex Flight Arena, demonstrating a 94.2% success rate under dynamic wind disturbances [3].

Anticipated Questions & Answers

This section details typical questions expected from the examination panel, alongside prepared answers that focus on defending key thesis contributions.

Question: Why did you choose deep reinforcement learning over classic Model Predictive Control for real-time trajectory generation?

While Model Predictive Control (MPC) provides robust safety guarantees through online optimization, it suffers from two major limitations in our application:

1. **Computational Latency:** The optimization problem becomes non-convex in the presence of arbitrary dynamic obstacles, leading to execution latencies exceeding 150 ms.
2. **Sensor Integration:** MPC cannot directly process high-dimensional spatial sensor data (like LiDAR point clouds) without substantial preprocessing overhead.

By utilizing **SynapseNet**, we shift the computational burden to offline training. During online deployment, inference requires only 11.4 ms, enabling fast feedback loops. Safety is maintained by deploying the **Apex** Safety Filter, which runs as a lightweight supervisory layer.

Question: How do you guarantee that the safety filter will not lead to deadlocks in highly cluttered environments?

The **Apex** Safety Filter operates by projecting the nominal DRL control input onto a set of safe control inputs defined by Control Barrier Functions (CBFs). In dense clutter, this projection can result in a null control space or local minima (deadlocks). We resolve this using a two-stage approach:

1. **State Expansion:** When the safe control set becomes empty, the controller inflates safety boundaries and executes a recovery sequence that redirects the drone along its historical trajectory.
2. **Active Sensing:** The agent performs small rotational maneuvers to gather missing spatial data in blind spots, resolving occlusions that cause false-positive local minima.

This ensures that the UAV retreats safely rather than becoming trapped or colliding.

Question: Your thesis assumes bounded sensor noise. How does the system handle complete sensor dropout or hardware failure?

Under extreme sensor dropouts, the assumption of bounded noise is violated. In this case:

- The epistemic uncertainty estimation module of **SynapseNet** detects a massive spike in out-of-distribution indicators, causing the uncertainty metric to exceed our critical threshold ($\theta_{max} = 0.85$).
- The **Apex** Safety Filter immediately overrides the navigation system and triggers a fallback protocol: the quadcopter transitions to an autonomous hover state, utilizing low-level optical flow and inertial sensors to perform a controlled descent.

This fail-safe mechanism ensures the UAV lands safely even in the event of primary sensor failure.

Defence Strategy & Weak-points Table

A critical aspect of the defense is acknowledging potential limitations and presenting a structured response that references the thesis content. Table 4.1 lists the identified weaknesses, potential criticisms from examiners, and the planned defense strategy.

Table 4.1. Strategic Defense Matrix

Weakness Area	Anticipated Criticism	Defense & Reference
Sim-to-Real Gap	“The safety filter was evaluated primarily in the Nimbus simulator. Real-world turbulence may violate safety bounds.”	The system was validated physically at the Apex Flight Arena under wind speeds of 5 m/s. The simulation models were calibrated using real wind-gust data to close the gap. See Section 5.3 (page 112) [2].
Sensor Constraints	“Solid-state LiDAR sensors have limited field-of-view, making the UAV blind to obstacles from the rear.”	The temporal convolution network captures the history of obstacle locations, building a localized memory map that keeps track of rear obstacles during forward motion. See Chapter 3 [1].
Computational Cost	“The Jetson Orin Nano may run out of memory when deploying a multi-agent swarm controller.”	Quantization using TensorRT reduces the model memory footprint by 68% while maintaining a latency of 11.4 ms, which falls within the required limits. See Section 5.4 [3].

4.1 Summary of Self-Criticism

While the experimental results validate our core claims, future iterations should focus on:

- **Decentralized Communication:** Currently, the swarm relies on localized sensing. Integrating decentralized ad-hoc communication networks could allow cooperative path planning.

- **Heterogeneous Swarms:** Testing the compatibility of **SynapseNet** across quadcopters with differing payload capacities and thrust-to-weight ratios.

References

- [1] Sarah L. Jenkins and Alexander Mercer. “Spatio-Temporal Path Planning in GPS-Denied Environments using Attention Networks”. In: *Meridian Journal of Artificial Intelligence* 8.1 (2026), pp. 89–104.
- [2] Alexander Mercer and Marcus Vance. “Epistemic Uncertainty and Sim-to-Real Transfer in Safety-Critical Robotics”. In: *Journal of Autonomous Robotic Systems* 14.2 (2025), pp. 112–128.
- [3] Elena Rostova, Alexander Mercer, and Aline Dubois. “Real-Time Neural Inference Quantization on Resource-Constrained Edge Platforms”. In: *Proceedings of the Apex Robotics Conference (ARC)*. Apex Robotics Society. 2026, pp. 45–52.