

1. SYSC 4907 P1: NeuroLogic Shield: A Modular Robustness Assessment Platform for Neural and Logic-Based AI Models

Modern machine learning and AI systems, e.g., whether neural networks and logic-based models, can behave unpredictably when exposed to small, carefully crafted input perturbations. This project focuses on giving students hands-on experience with adversarial robustness, a critical skill for validating AI models before deployment. Students will build a compact “NeuroLogic Shield” platform that evaluates how two types of models: a simple multilayer perceptron (MLP) and a differentiable logic-guided network (DLGN), respond to a single, well-understood adversarial method: the Fast Gradient Sign Method (FGSM). By limiting the scope to one attack and two model classes, the project stays accessible while still offering deep value in understanding model behavior, debugging, and system-level evaluation.

The team will implement FGSM from first principles, integrate it into a clean and modular evaluation pipeline, and run controlled experiments to assess each model’s robustness. They will design lightweight visualizations and a simple reporting interface that clearly communicates attack impact, accuracy degradation, and model-specific vulnerabilities. The emphasis is on correctness, engineering clarity, and reproducible experimentation rather than building a full adversarial suite. With support in the form of starter code, background training, and close supervision, students can complete a technically meaningful project that teaches practical ML auditing skills and potentially produces a small open-source contribution or workshop-grade technical report.

2. SYSC 4907 P2: QuantumVerify: Design and Evaluation of a Post-Quantum Secure Document Signing Platform

As quantum computing advances, many widely used digital-signature schemes—such as RSA and elliptic-curve signatures—face long-term vulnerabilities, raising concerns about the authenticity of documents that must remain trustworthy for years or decades. This project explores that emerging security challenge by building “QuantumVerify,” a functional document-signing platform that relies entirely on standardized post-quantum signature algorithms. Students will integrate CRYSTALS-Dilithium and Falcon into a realistic signing workflow to examine how these quantum-resistant schemes behave in practice. The project blends security fundamentals with applied engineering: instead of studying cryptographic math, students gain hands-on experience with how post-quantum primitives affect file sizes, performance, user workflows, and the feasibility of long-term digital authenticity.

The team will develop a working document-signing and verification application, run controlled experiments on files of varying sizes, and benchmark key metrics such as signing time, verification latency, key sizes, and signature overhead. The focus is on reproducible engineering evaluation and clear reporting rather than building a full production system. Students will also deploy the platform to a Raspberry Pi or microcontroller to compare desktop-class and embedded-class performance, highlighting real-world constraints faced in IoT, compliance, and low-power environments. The final result offers both practical insight into the trade-offs between Dilithium and Falcon and a hands-on experience with technologies that organizations will increasingly rely on as part of the post-quantum transition.

3. SYSC 4907 P3: Privacy and Transparency in Modern AI Services

Modern AI services are becoming central to communication, business operations, and software platforms across Canada, yet most users still have little visibility into what happens to their data once it’s submitted to an AI tool. This project explores the emerging “AI Transparency Gap”: a growing uncertainty around who has access to user data, how long it’s stored, whether it contributes to future model training, and where in the world it travels. For organizations required to comply with Canadian privacy principles, especially

PIPEDA's Accountability and Cross-Border Transfer guidelines, this lack of clarity creates practical and legal risks. The goal of this project is to make those hidden processes visible in a way that is understandable to both technical and non-technical audiences.

The student team will investigate how major AI-as-a-Service platforms manage data by creating a structured taxonomy of privacy and access risks, then turning those findings into a public-facing transparency dashboard. Students will analyze documentation, APIs, and platform behaviour to classify services based on access hierarchy, data residency, training ingestion, cross-session memory, and output ownership. Using these insights, they will build a web-based "AI Privacy Pulse" tool that presents simple scorecards, like a nutrition label, for each AI provider. This tool will help Canadian users and small organizations make informed decisions about which AI services align with their privacy expectations.