CICI: UCSS: Secure Machine Learning as a Service for Collaborative Scientific Research

PI: Dr. Yushun Dong (Florida State University);

Co-PI: Dr. Mengxin Zheng (University of Central Florida); Dr. Neil Gong (Duke University)

Project URL: https://yushundong.github.io/CICI2026/

Collaboration and Usability Challenge Addressed

Thrust 1 — Model Confidentiality (against MEAs)

Attackers can steal ML models via API querying. We propose proactive defenses combining sensitivity-regularized training and decision-focused ensemble modules at deployment to make replication computationally infeasible while preserving model utility.

Thrust 2 — Data Secrecy (against MIAs)

Adversaries can reconstruct sensitive training data from model outputs. We develop a Privacy Funnel framework that suppresses high-risk input prototypes in representations, alongside adaptive noise injection targeting suspicious queries to protect data without degrading performance.

Thrust 3 — Inference Privacy (against data breaches)

User queries and outputs in CI environments are vulnerable to interception. We design efficient homomorphic encryption with ciphertext-aware pruning and polynomial reduction to enable practical, fully encrypted inference.

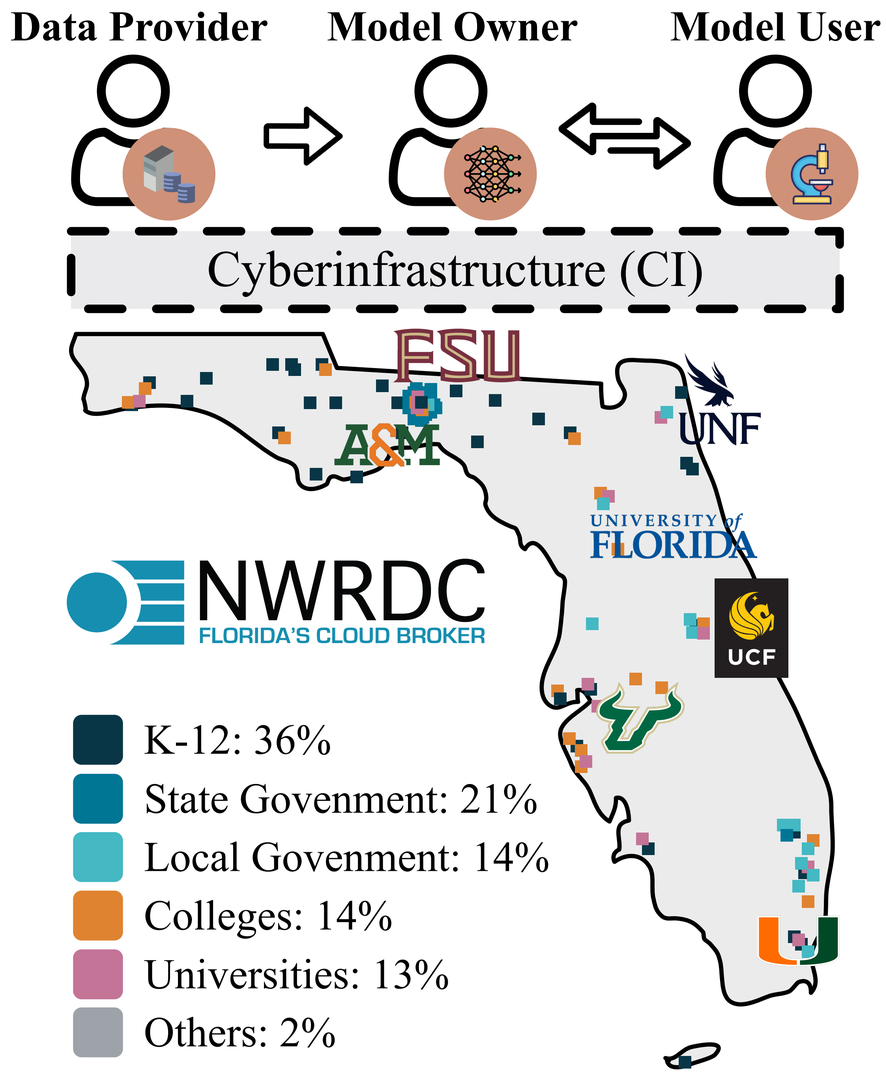

Technical Cybersecurity Solution

- Model Confidentiality: Sensitivity Regularization; Model Ensemble.
- Data Secrecy: Privacy Funnel; Suspicious Input Detection & Obfuscation.
- Inference Privacy: Efficient Encryption; Dynamic Polynomial Reduction.

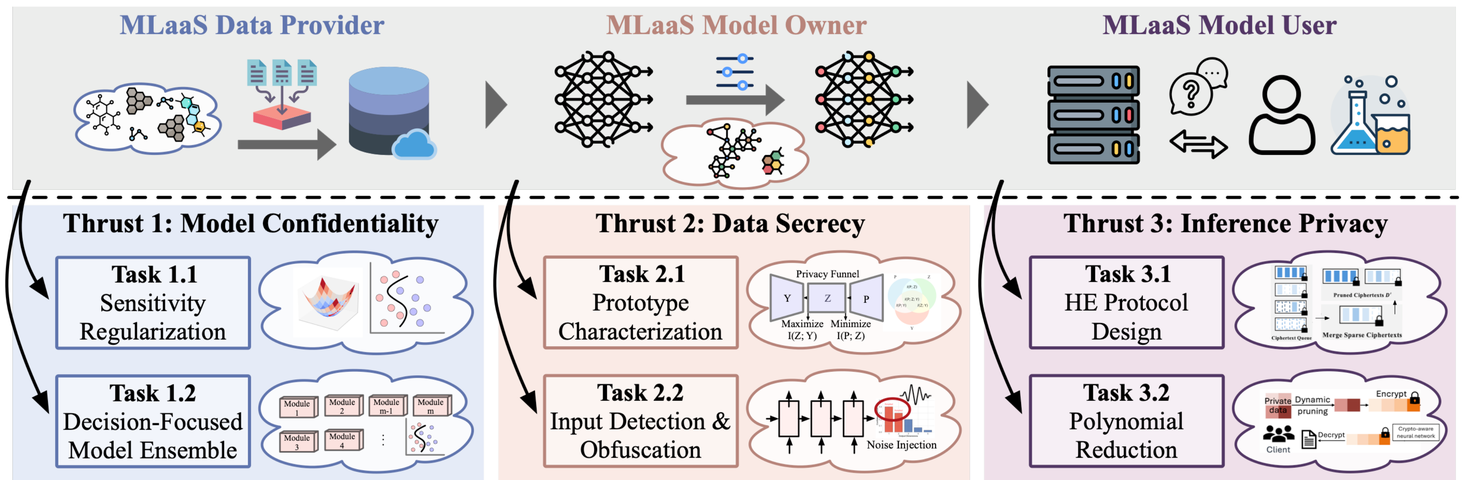

Benefits to Scientific Cyberinfrastructure

- Strengthens security, integrity, and reproducibility of MLaaS-based scientific collaboration workflows by mitigating model, data, and inference-level vulnerabilities.
- Enhances trust in collaborative CI for sensitive domains (healthcare, disaster response, genomics) by enabling secure and reliable model access and data sharing across institutions.

Result Dissemination Plans

- Stolen model fidelity↓; reconstructed feature distance↑; encrypted inference latency↓; negligible task-relevant utility loss↓.
- Open-source software; Public datasets & benchmarks; Tutorials, user guides, workshops, etc.
- Audience: CI admins, domain scientists, research institutions.

Risks Versus Potential For Advances

- **Risks**: Performance overhead from encryption; resistance to adoption if workflows disrupted;
- **Mitigation**: Efficiency-focused design; modular integration; continuous usability feedback loops.
- **Payoff**: A deployable, modular security framework that protects models, data, and inference in collaborative MLaaS environments.