ΠΑΝΕΠΙΣΤΗΜΙΟ ΠΕΙΡΑΙΩΣ
ΤΜΗΜΑ ΨΗΦΙΑΚΩΝ ΣΥΣΤΗΜΑΤΩΝ

ΠΡΟΓΡΑΜΜΑ ΜΕΤΑΠΤΥΧΙΑΚΩΝ ΣΠΟΥΔΩΝ
ΚΥΒΕΡΝΟΑΣΦΑΛΕΙΑ ΚΑΙ ΤΕΧΝΟΛΟΓΙΕΣ ΤΕΧΝΗΤΗΣ ΝΟΗΜΟΣΥΝΗΣ

ΘΕΜΑΤΑ ΔΙΠΛΩΜΑΤΙΚΩΝ ΕΡΓΑΣΙΩΝ

(ΑΚΑΔΗΜΑΪΚΟ ΕΤΟΣ 2024-25)

ΕΠΙΒΛΕΠΩΝ/ΟΥΣΑ:

Α/Α	Θέμα Διπλωματικής Εργασίας	Σύντομη Περιγραφή [1]	Επιβλέπων
1.	**Adversarial attacks and Mitigation Techniques against Few Shot Learning Models**	The purpose of this dissertation is to perform a survey on adversarial attacks (such as data poisoning) that can be applied against Few Shot Learning models and present mitigation techniques against them. (MB)	
2.	**A Survey of Adversarial Attacks in Transfer Learning**	Create a taxonomy of adversarial attacks specifically targeting transfer learning and evaluate proposed defense mechanisms. (MB)	
3.	**Backdoor Attacks in Transfer Learning: A Survey of Methods, Targets, and Defenses**	The purpose of this dissertation is to perform a survey on backdoor attacks (e.g., trojan triggers embedded during pretraining) and examine the persistence of attacks after	

		transfer learning. Moreover, the dissertation will include an evaluation of the defense mechanisms against the identified attacks. (MB)	
4.	**Cybersecurity Data Fabric**	To design, prototype, and evaluate a secure, decentralized data architecture leveraging blockchain and decentralized storage (Hyperledger Fabric & IPFS or EBSI), aimed at ensuring data integrity, provenance, and fine-grained access control in multi-organizational environments. Expectations: ● A working **prototype** of a decentralized cybersecurity data fabric. ● Access control (Self-sovereign-identity) ● Smart contracts for audit logging, and data lifecycle governance. ● Evaluation of **performance, security**, and **resilience** of the solution. ● A conceptual model for applying this fabric in real-world domains (e.g., healthcare, supply chains, public sector). ● Use case scenario (AF)	

5.	**Evaluating IoT Device Security Through Penetration Testing with Flipper Zero: Methods, Exploits, and Countermeasures**	To assess the security posture of IoT devices by conducting systematic penetration tests using Flipper Zero, identifying common vulnerabilities in wireless communication protocols, and proposing practical mitigation strategies. (AF)	
6.	**Security Assessment of Open FPGA**	To investigate the security posture of OpenFPGA, identify potential vulnerabilities in the hardware architecture and toolchains, and propose countermeasures or secure development practices. **Expectations:** ● A **threat model** tailored to open-source FPGA workflows. ● Identification and classification of vulnerabilities across the open FPGA stack. ● A security evaluation of OpenFPGA. Guidelines for securing the open FPGA development and deployment process. (AF)	

7.	**Zero Trust for LLM Agents: Detection and Mitigation of Adversarial Behaviors in Multi-Agent AI Systems**	To propose a Zero Trust-inspired framework to detect adversarial behavior between AI agents, and restrict or sandbox agents based on trust scores, policies, and behavior. Expectations: 1. multi-agent simulation using LLM-based agents with tasks (i.e., data analysis, report writing, email filtering) connected via an agent framework (i.e., autogen). 2. Adversarial Attack execution 3. Application of detection & mitigation mechanisms 4. Evaluation of attacks and detections and mitigation later on. (AF)	
8.	**Self-Sovereign Identity for LLM-Agent Access Control in Multi-Agent Systems**	Combine SSI with agentic AI to create identity-bound permissions for autonomous LLM agents (like AutoGen, CrewAI, LangGraph agents). Each agent will carry a verifiable credential (VC) defining who it is, what it can access, and for how long. Evaluation of credential issuance/verification time, agent misuse prevention rate, latency, scalability, revocation propagation delay. (AF)	

9.	**AI Powered AV for Medical IoT Malware Detection**	This involves detecting Linux malware targeted at medical IoT devices, utilizing deep learning networks to analyze malware which has been converted into image files. Both grayscale and RGB images will be considered and a cost-benefit analysis will decide which images will be selected. This also includes evaluating model performance using RNNs, CNNs, and autoencoders, and assessing robustness against adversarial attacks generated by a GAN, while also addressing memory challenges in processing large image files. Numerical results for the malware classification will be included. (AF)	
10.	**Beyond Shamir: Adaptive Secret Sharing for Dynamic, Decentralized Environments**	Implement and evaluate adaptive secret sharing schemes where participants can join/leave without resharing from scratch. Explore proactive sharing, share refreshing, and mobile adversaries. **Expected Outcome:** Python/Rust implementation of an adaptive scheme (e.g., proactive Shamir or threshold ramp schemes), simulation with dynamic group changes, performance metrics. (AV)	

11.	**Verifiable Secret Sharing in Adversarial MPC: Towards Practicality with Lattice Assumptions**	Prototype lattice-based VSS using Ring-LWE or LWE, exploring commitments and proofs (e.g., ZK proofs for share validity) and analyze resistance to quantum attacks. **Expected Outcome:** A Rust or Python proof-of-concept of lattice-based VSS using libraries like concrete,TFHE (AV)	
12.	**From Pedersen to Aggregatable DKG: A Survey and Implementation Benchmark of Modern DKG Protocols**	Survey and re-implement multiple DKG protocols (Pedersen, Joint-Feldman, Scrape, FROST-DKG). Benchmark communication rounds, message sizes, and fault tolerance. **Expected Outcome:** Python or Rust modular implementations of DKG variants. Graphs showing efficiency and fault tolerance comparisons under simulated networks. (AV)	

13.	**VSS to the Edge: Lightweight Verifiable Secret Sharing for IoT and 6G Networks**	Design and test VSS protocols suitable for resource-constrained devices (Raspberry Pi or other IoT devices). Optimize for bandwidth, CPU, and memory usage. **Expected Outcome:** Lightweight VSS protocol (Pedersen or novel), deployed on emulated IoT nodes. Performance report under bandwidth and latency constraints. (AV)	
14.	**Round-Efficient Threshold ECDSA and its Post-Quantum Successors**	Implement threshold ECDSA (e.g., GG18, FROST) and explore adaptations using PQ-safe primitives like Dilithium or Picnic or even MAYO. Compare round complexity and signature size. **Expected Outcome:** Working ECDSA threshold implementation (Python or Rust), plus a PQ prototype. Experiments on rounds, signature latency, and fault tolerance. (AV)	

15.	**Authenticated KEMs in the Post-Quantum World: Lattice-Based Constructions and Limitations**	Implement Authenticated KEM constructions using Kyber/Frodo combined with digital signatures or one-pass AKE protocols. Analyze security trade-offs and efficiency. **Expected Outcome:** Python/Rust implementation of authenticated KEM variants (e.g., signed-KEM or tag-based encryption), tested in simulated client-server communication. (AV)	
16.	**From Kyber to KEMTLS: Integrating PQC KEMs in Real-World Authenticated Key Exchange Protocols**	Study and implement KEMTLS-like protocols using Kyber/Frodo and hybrid mechanisms with classical TLS. Perform handshake latency and bandwidth analysis. **Expected Outcome:** Experimental Python/TLS code (e.g., modified tslite-ng or liboqs or "custom handshake logic") simulating PQ KEMTLS handshakes. Plots of handshake time, message sizes. (AV)	
17.	**Sign of the Times: What's Next After NIST's PQ Signature Round – A Comparative Landscape**	Survey and implement leading PQ signature schemes (Dilithium, SPHINCS+, Falcon). Compare size, signing/verifying speed, and security assumptions. Survey	

		the new proposals like FAEST,HAWK,MAYO,SQISign, QR-UOV etc. **Expected Outcome:** Benchmarking suite in Python (using NIST reference code bindings or wrappers). Side-by-side performance analysis under varying message sizes and platforms. (AV)	
18.	**Zero-Knowledge and Secret Sharing: Building Verifiable Computation from Distributed Secrets**	Combine ZK proofs (e.g., Schnorr or Bulletproofs) with secret-shared data to verify correctness of computations without revealing inputs. **Expected Outcome:** Prototype using Python or Rust (e.g., zk-intercafe, circom, or bulletproofs in Rust) to demonstrate secret-shared voting or sum verification with ZK proofs. (AV)	
19.	**Threshold Cryptography Meets Federated Learning: Secure Aggregation with VSS and DKG**	Implement secure aggregation in federated learning using VSS + DKG to compute global models without revealing local weights. Protect against dropouts. Can even adapt FHE. (AV)	

20.	**Lattice-Based Threshold Primitives: Building the Future of Secure Multi-Party Systems**	Design and prototype threshold key generation and signing with lattice-based schemes (Kyber, Dilithium). Explore homomorphic properties of RLWE-based schemes. **Expected Outcome:** Experimental Rust or Python code implementing lattice-based threshold signing or encryption. Evaluate correctness, error rates, and noise growth. (AV)	
21.	**Secure AI Model Certification Pipeline under the EU AI Act: A Technical Framework**	Investigate and prototype a security-focused AI assurance pipeline for high-risk AI systems, aligned with the EU AI Act and NIS2. Emphasis on model documentation, robustness evaluation, logging, traceability, and post-deployment monitoring. Use case: AI for critical infrastructure (e.g., transport or healthcare). **Expected Outcome:** Taxonomy of technical and regulatory requirements; prototype of a traceability framework (e.g., model lineage with secure logs); evaluation on simulated audit scenarios. (VB)	
22.	**keep Cross-Domain Identity Federation for AI Agents using Blockchain and SSI**	Propose and implement a cross-domain identity federation framework for autonomous AI agents using Self-Sovereign Identity (SSI) anchored on blockchain.	

		Explore interoperability across multiple blockchains and credential issuers (e.g., EU eIDAS 2.0). **Expected Outcome:** Working SSI-based identity management prototype; integration with LLM agents; benchmarking identity verification latency, scalability, and revocation. (VB)	
23.	**AI for Threat Detection and Incident Response in Critical Infrastructures**	This topic explores the design and deployment of AI-based models for detecting cyber threats and automating incident response in real-time across critical infrastructure environments. **Expected Outcome:** Develop, train, and evaluate AI-driven threat detection systems tailored to critical infrastructure scenarios, with a deep understanding of data pipelines, adversarial robustness, and deployment challenges. (GC)	
24.	**Privacy-Preserving Machine Learning and Federated AI**	The topic introduces advanced techniques in privacy-preserving machine learning, including federated learning and secure multi-party computation, applied to distributed AI environments such as edge and IoT systems. **Expected Outcome:**	

		Hands-on experience designing privacy-aware AI system that comply with regulations (e.g., GDPR), enabling decentralized training and inference without compromising sensitive data. (GC)	
25.	**AI-Based Cyber Risk Assessment and Governance**	This topic focuses on applying AI to automate cyber risk analysis, decision-making, and governance processes in alignment with cybersecurity standards and regulatory frameworks. **Expected Outcome:** Build AI-assisted cyber risk models, simulate attack scenarios, and support compliance with national and international cybersecurity policies in complex ICT ecosystems. (GC)	
26.	**Development of Attack Defence CTF Services**	Capture the Flag (CTF) competitions may have different styles and modes. The Attack/Defense type of competition is a CTF where multiple services are given to the teams and they are tasked with hacking the services of the other teams while protecting their own. The purpose of this dissertation is to develop 3 Attack-Defense services that could be used through CTF competitions to train users. Suggested to have: Intermediate experience in CTF competitions, Some experience in A/D competitions	

		Expected Outcome: Investigation into how A/D services work, and development of 4 services (AG)	
27.	**AI assisted CTF challenge development**	Investigate ways AI can be used to assist in the development of CTF challenges (e.g. new Image generation, code generation). Design, test and evaluate methodologies to develop configurable CTF challenges that can be mutated using generative AI. **Expected Outcome:** A methodology utialising AI systems for assing in CTF challenges creation. (AG)	
28.	**Development of gamified cybersecurity trainings**	Design and develop cybersecurity training modules in the form of interactive games for the Hack'n'Lean platform, utilising emulated mock environments and minigames to mix theory with practice in a game-like approach. **Expected Outcome:** 1 series of interactive modules (4-8 modules) V	

[1] *Στη σύντομη περιγραφή της διπλωματικής εργασίας θα πρέπει να αναφέρεται σαφώς το τελικό αποτέλεσμα αυτής (Εκπόνηση Μελέτης, Ανάπτυξη Συστήματος / Υπηρεσίας, Ανασκόπηση / Συγκριτική Αξιολόγηση, κτλ).*

ΠΑΝΕΠΙΣΤΗΜΙΟ ΠΕΙΡΑΙΩΣ
ΤΜΗΜΑ ΨΗΦΙΑΚΩΝ ΣΥΣΤΗΜΑΤΩΝ

ΠΡΟΓΡΑΜΜΑ ΜΕΤΑΠΤΥΧΙΑΚΩΝ ΣΠΟΥΔΩΝ
ΚΥΒΕΡΝΟΑΣΦΑΛΕΙΑ ΚΑΙ ΤΕΧΝΟΛΟΓΙΕΣ ΤΕΧΝΗΤΗΣ ΝΟΗΜΟΣΥΝΗΣ