

KUDELSKI
—★— **LABS**

KUDELSKI
SECURITY 

KUDELSKI
IOT THINGS

Kudelski Group

INTERNSHIPS 2027

TABLE OF CONTENTS

[AST-LAB-01] Drive the hyperparameters research for Attacks models	4
[AST-SEC-01] Side Channel and Fault Attack Investigations on Hamming Quasi-Cyclic (HQC) Post-Quantum Cryptography	5
[AST-SEC-02] FrodoKEM: Algorithm, Proof of Concept and Constrained-Device Feasibility	6
[AST-SEC-03] Classic McEliece: Algorithm, Proof of Concept and Physical Security State of the Art.....	8
[AST-SEC-04] Quantum Random Number Generation: State of the art and certification	10
[AST-SEC-05] AI-Assisted Cryptanalysis: Neural Distinguishers and Learning-Based Attacks on Lattices.....	12
[AST-SEC-06] Hybrid Post-Quantum Cryptography and Quantum Key Distribution Architecture	14
[AST-AI-01] State-of-the-Art Exploration of AI World Models	16
[AST-AI-02] State-of-the-Art Exploration of Audio Foundation Models for Detection and Situational Awareness	18
[AST-AI-03] Synthetic World Simulation Platform.....	20
[AST-AI-04] Context-Aware Robot using Local LLMs.....	22
[COE-AI-01] Agentic Web Monitoring for Audiovisual Piracy	24
[KS-MDR-01] MDR business operations & technology EPFL master's intern.....	25

HOW TO APPLY

- Open positions are advertised on our website careers.nagra.com
 - All applications must be done through our careers page
 - Trainee – **Ref 15845**
 - Mention the internship number(s) in motivation letter e.g. **[AST-LAB-01]**
-

[AST-LAB-01] Drive the hyperparameters research for Attacks models

Topics: **AI- Model Architecture-Hyper parameters**

Master Thesis Possible: **Yes**

Development of a solution to reduce the overhead induced by hyperparameter search for our side-channel or fault attack models against secure devices.

Details

Today, the recent development of AI domain leads to an explosion of models, all accompanied by an exponential growth in the number of hyperparameters. In the context of our physical attack activities (attacks through side channels, fault attacks), it becomes impossible to select the best solution and find adequate parameters in a reasonable time for our learnings.

The purpose of the internship is to identify the best approach to provide a solution to this problem. Two options can be considered:

- The first could consist of using an agent to find the adequate hyper parameters by monitoring GPU training (done in Matlab/Pytorch) as shown in [1].
- The second potential solution aims to regroup all the models' trials in one global model with stacking/bagging approach [2].

Key Objectives

- Advance in the development of a method to find the best model/hyperparameters suitable to our side channel attacks.
- Test and Validate efficiency on concrete attacks dataset
- Provide a clear documentation support for the use of tools

Technical Stack:

- Matlab/Python/C
- Good understanding of the deep learning architectures (MLP, CNN, Transformer,)
- Good mastery of software environment (Linux, Windows, WSL, JupyterHub,...)
- Interest to GPU hardware
- Understanding Side channel domain is a plus

Expected Outcomes:

- A robust and easy to implement solution for our daily attack work
- Potential extension for other domains of application.

Team & Collaboration:

- Onboarded within the Kudelski security Lab

[1] <https://github.com/karpathy/autoresearch>

[2] Deep Stacking Ensemble Learning Applied to Profiling Side-Channel Attacks

[AST-SEC-01] Side Channel and Fault Attack Investigations on Hamming Quasi-Cyclic (HQC) Post-Quantum Cryptography

Topics: **Quantum security – Side Channel- Fault Attacks** Master Thesis Possible: **Yes**

This subject aims to evaluate the physical security of HQC implementations against side channel attacks (SCA) and fault injection attacks (FIA). The goal is to identify potential vulnerabilities in embedded or hardware-based deployments and propose effective countermeasures.

Details

HQC is a code-based post-quantum cryptographic scheme relying on the hardness of decoding problems in Hamming metric and quasi-cyclic codes. While mathematically secure against quantum adversaries, its real-world implementations may leak sensitive information through physical channels or be disrupted by induced faults. Side channel and fault attacks have proven effective against many key encapsulation cryptographic schemes, including other post-quantum candidates like ML-KEM and BIKE. This thesis will explore how HQC stands up to such threats and how to harden its implementations.

Key Objectives

- Identify what are the main side channel leakage vectors in HQC implementations
- Identify how vulnerable is HQC to fault injection attacks
- Define which countermeasures can be applied to protect HQC implementations
- Compare HQC to other PQ schemes in terms of physical security and its securitization.

Technical Stack:

- HQC Reference Implementation (e.g., PQCclean)
- Python / C for simulation and scripting
- ChipWhisperer (optional) for SCA and FIA
- Python libraries: NumPy, SciPy, Matplotlib

Expected Outcomes:

- A security assessment of HQC implementations under physical attack scenarios
- Simulated or experimental attack results
- Comparative analysis with other PQC schemes
- Recommendations for secure HQC deployment.

Team & Collaboration:

- Onboarded within the Advanced Research team, collaborating with cybersecurity, Lab attacks specialists.
- Potential for cross-functional integration with other secure communication initiatives.

[AST-SEC-02] FrodoKEM: Algorithm, Proof of Concept and Constrained-Device Feasibility

Topics: **Quantum security – Side Channel- Fault Attacks** Master Thesis Possible: **Yes**

This subject aims to evaluate the physical security of FrodoKEM implementations against side channel attacks (SCA) and fault injection attacks (FIA). The goal is to identify potential vulnerabilities in embedded or hardware-based deployments and propose effective countermeasures.

Details

FrodoKEM is a plain-LWE key encapsulation mechanism, without any ring or module structure. This conservative design is the reason why it is recommended by BSI and ANSSI, but it also makes it significantly heavier than ML-KEM in terms of bandwidth and matrix arithmetic. Recent work has shown message-recovery and key-recovery attacks exploiting ciphertext malleability and the Fujisaki–Okamoto re-encryption step, including first-order attacks against a first-order masked hardware implementation. This subject focuses on making FrodoKEM work well and safely on constrained targets, with the protected variants in scope from the start.

Key Objectives

- Implement and benchmark FrodoKEM on an embedded target (RISC-V)
- Identify what are the main bottlenecks (matrix generation, memory footprint, bandwidth)
- Evaluate the cost of masking and shuffling countermeasures on performance
- Define which hybrid constructions with ML-KEM are relevant for European deployments
- Compare FrodoKEM to ML-KEM in terms of performance, bandwidth and physical security.

Technical Stack:

- FrodoKEM Reference Implementation
- Python / C for simulation and scripting
- Chip Whisperer (optional) for SCA and FIA
- Python libraries: NumPy, SciPy, Matplotlib

Expected Outcomes:

- A working proof of concept on a constrained platform
- Performance, memory and bandwidth benchmark results
- Comparative analysis with other lattice-based KEMs
- Recommendations for FrodoKEM deployment and for hybrid PQC modes

Team & Collaboration:

- Onboarded within the Advanced Research team, collaborating with cybersecurity, Lab attacks specialists.
- Potential for cross-functional integration with other secure communication initiatives.

Reference Papers

1. *Countermeasures against side-channel attacks in FrodoKEM: an analysis*, Journal of Cryptographic Engineering, 2026 [\[link.springer.com\]](https://link.springer.com)
2. *In-depth Analysis of a Side-Channel Message Recovery Attack Against FrodoKEM*, IEEE, 2025 [\[ieeexplore.ieee.org\]](https://ieeexplore.ieee.org)
3. *The Fellowship of the Leak: Power Analysis of a Masked FrodoKEM Implementation*, IEEE, 2025 [\[ieeexplore.ieee.org\]](https://ieeexplore.ieee.org)
4. Y. Tanaka, R. Ueno, K. Xagawa, A. Ito, J. Takahashi, N. Homma, *Multiple-Valued Plaintext-Checking Side-Channel Attacks on Post-Quantum KEMs*, TCHES 2023(3) [\[tches.iacr.org\]](https://tches.iacr.org)
5. J. W. Bos et al., *Assessment of side-channel leakage on FrodoKEM — as modelled in the leaky-LWE-Estimator framework* [\[deepwiki.com\]](https://deepwiki.com)
6. P.-A. Berthet, *Side channel resistance of Europe's choices for lattice-based post-quantum cryptography: FALCON and FrodoKEM*, PhD thesis, 2025 [\[theses.hal.science\]](https://theses.hal.science)

[AST-SEC-03] Classic McEliece: Algorithm, Proof of Concept and Physical Security State of the Art

Topics: **Quantum security – Side Channel- Fault Attacks** Master Thesis Possible: **Yes**

This subject aims to build a working understanding and proof of concept of Classic McEliece, and to establish a consolidated state of the art of its resistance to side channel and fault injection attacks. The goal is to characterize its attack surface and assess the real cost of protecting it.

Details

Classic McEliece is a code-based KEM whose security relies on the hardness of decoding binary Goppa codes, and which benefits from more than forty years of cryptanalytic scrutiny. It is a strong candidate for long-term and sovereign use cases despite its very large public keys. Key-recovery side channel attacks have been demonstrated on both FPGA and Cortex-M4 implementations, exploiting leakage in the additive FFT used to evaluate the error locator polynomial, and more recently in the syndrome computation of the reference decoder. The first protected hardware implementation reports an area overhead of 4× to 7×, which makes the cost/benefit analysis a central question for embedded deployment.

Key Objectives

- Implement and benchmark McEliece on an embedded target (RISC-V)
- Identify what are the main side channel leakage vectors in Classic McEliece
- Identify how vulnerable Classic McEliece is to fault injection attacks
- Define which countermeasures can be applied and quantify their latency and memory cost
- Compare Classic McEliece to HQC and ML-KEM in terms of physical security and its saucerization
- Assess the practical feasibility of Classic McEliece on a constrained secure element.

Technical Stack:

- Classic McEliece reference implementation and FPGA reference design
- Python / C for simulation and scripting; SageMath for code-based arithmetic
- ChipWhisperer for SCA and FIA
- Python libraries: NumPy, SciPy, Matplotlib, scikit-learn

Expected Outcomes:

- A consolidated state of the art and a security assessment of Classic McEliece implementations
- Simulated or experimental attack results
- Comparative analysis with other code-based PQC schemes

- Recommendations for secure Classic McEliece deployment

Team & Collaboration:

- Onboarded within the Advanced Research team, collaborating with cybersecurity, Lab attacks specialists.
- Potential for cross-functional integration with other secure communication initiatives.

Reference Papers

1. Q. Guo, A. Johansson, T. Johansson, *A Key-Recovery Side-Channel Attack on Classic McEliece Implementations*, TCHES 2022(4) [\[tches.iacr.org\]](https://tches.iacr.org/)
2. M. Bulois, P.-L. Cayrel, V.-F. Drăgoi, V. Grosso, *Algebraic Key-Recovery Side-Channel Attack on Classic McEliece*, SAC 2025 [\[sacworkshop.org\]](https://sacworkshop.org/)
3. P. Gan, P. Ravi, K. Raj, A. Baksi, A. Chattopadhyay, *Classic McEliece Hardware Implementation with Enhanced Side-Channel and Fault Resistance*, 2024 [\[techrxiv.org\]](https://techrxiv.org/)
4. Classic McEliece team, *Classic McEliece: Papers* — curated bibliography on ISD and implementation attacks [\[classic.mceliece.org\]](https://classic.mceliece.org/)
5. Classic McEliece, *NIST PQC Round 4 specification and design rationale*

[AST-SEC-04] Quantum Random Number Generation: State of the art and certification

Topics: **Quantum security – Randomness**

Master Thesis Possible: **Yes**

This subject aims to establish a state of the art of quantum random number generator (QRNG) architectures and to define how such generators can be evaluated and certified. The goal is to produce an evaluation methodology comparable to the one used for classical physical TRNGs.

Details

QRNGs exploit quantum phenomena such as photon arrival times, vacuum fluctuations or photon path splitting to produce entropy whose unpredictability is grounded in physics rather than in complexity assumption. The field spans fully trusted-device designs, semi-device-independent protocols and device-independent protocols based on Bell or Leggett–Garg inequality violation, each with very different assumptions and very different practical cost. From a certification perspective the central question is entropy estimation: a QRNG must justify a min-entropy rate under an explicit stochastic model, exactly as required by AIS 20/31 and SP 800-90B for classical TRNGs. This subject directly extends the ring-oscillator and jitter-based entropy work already ongoing.

Key Objectives

- Survey QRNG architectures and classify them by entropy source and trust assumptions
- Identify how quantum entropy is estimated and what stochastic models are used
- Compare QRNG with classical physical TRNGs in terms of entropy rate, throughput and assurance
- Define which requirements of AIS 20/31 and SP 800-90B a QRNG can and cannot satisfy today
- Propose an evaluation methodology and a reference test plan

Technical Stack:

- Python for statistical analysis and entropy estimation
- NIST SP 800-22 / SP 800-90B test suites, AIS 31 test procedures
- Qiskit or equivalent for protocol simulation
- Python libraries: NumPy, SciPy, Matplotlib

Expected Outcomes:

- A comparative state of the art of QRNG technologies
- A documented entropy evaluation and certification methodology
- A mapping of QRNG properties to AIS 20/31 and SP 800-90B requirements
- Recommendations on the relevance of QRNG for KSE-class products

Team & Collaboration:

- Onboarded within the Advanced Research team, collaborating with cybersecurity, Lab attacks specialists.
- Potential for cross-functional integration with other secure communication initiatives.

Reference Papers

1. P. P. Nath, A. Sinha, U. Sinha, *Certified Random Number Generation using Quantum Computers*, Frontiers in Quantum Science and Technology, 2025 [\[frontiersin.org\]](https://frontiersin.org)
2. R. Okula, P. Mironowicz, *Extensive search of Shannon entropy-based randomness certification protocols*, arXiv:2511.22771 [\[arxiv.org\]](https://arxiv.org)
3. S. Mazzucchi, N. Leone, S. Azzini, L. Pavesi, V. Moretti, *Entropy certification of a realistic quantum random-number generator based on single-particle entanglement*, Phys. Rev. A 104, 022416 [\[link.aps.org\]](https://link.aps.org)
4. K. Akhil, K. Kanagalakshmi, *Beyond Ideal Randomness: A Comprehensive Study on Practical Quantum Random Number Generators for Cryptographic Security*, 2025 [\[ijcrt.org\]](https://ijcrt.org)
5. NIST SP 800-90B and BSI AIS 20/31 — entropy source validation requirements

[AST-SEC-05] AI-Assisted Cryptanalysis: Neural Distinguishers and Learning-Based Attacks on Lattices

Topics: **Quantum security – Cryptanalysis – AI**

Master Thesis Possible: **Yes**

This subject aims to establish, with reproducible evidence, where machine learning genuinely helps mathematical cryptanalysis and where it does not.

Details

Since Gohr's neural distinguisher against round-reduced Speck32/64, machine learning has become a legitimate cryptanalytic instrument, and transformer models have been applied to recovering sparse LWE secrets under reduced parameters. None of these results threaten a standardised parameter set. Establishing where the boundary lies is the objective.

Key Objectives

- Reproduce a neural distinguisher on a round-reduced symmetric primitive and quantify the gain over classical differential cryptanalysis
- Determine what the learned distinguisher captures, and whether it is interpretable as a classical property
- Evaluate learning-based attacks on sparse-secret LWE and establish the parameter regime where they stop working
- Document the gap to standardised parameters in orders of magnitude

(only if previous are completed by M5)

- Apply the methodology to an internal Kudelski primitive
- Prototype a source-traceable crypto-analysis assistant

Phase	Focus	Milestone
M1	State of the art; reproduce Gohr-style baseline	Baseline reproduced
M2	Extension and interpretability analysis	Gain quantified vs classical
M3–M4	Learning-based attacks on sparse-secret LWE	Parameter boundary established
M5	Critical assessment note: what is and is not threatened	Note reviewed
M6	Consolidation, report, final presentation; stretch items if time allows	Report + presentation

Technical Stack

- PyTorch; scikit-learn

-
- SageMath, fpylll, the Lattice Estimator for classical baselines
 - Python: NumPy, SciPy, Matplotlib

Expected Outcomes

- Reproducible neural distinguisher results with quantified gain and interpretability analysis
- A parameter-regime map for learning-based lattice attacks
- A defensible internal position note on the real cryptanalytic capability of AI
- Publishable material if the interpretability or boundary results prove novel

Team & Collaboration

- Advanced Research team, collaborating with cryptography and AI specialists
- Cross-functional link to the PQC migration advisory activity

[AST-SEC-06] Hybrid Post-Quantum Cryptography and Quantum Key Distribution Architecture

Topics: **Quantum security – Secure communication**

Master Thesis Possible: **Yes**

This subject aims to design and evaluate architecture combining QKD and post-quantum cryptography. The goal is to determine how the two technologies complement each other, in particular regarding authentication, which QKD does not address.

Details

The existing QKD subject already notes that authentication of the parties is not solved by QKD itself. This is the central practical issue: a QKD link must be authenticated by a classical mechanism, and in a post-quantum world that mechanism itself must be quantum resistant. A hybrid architecture therefore uses PQC for authentication and key establishment and QKD for key agreement over the optical layer, with a key combiner ensuring that the resulting key is at least as strong as the strongest component. National agencies differ significantly in their positions on QKD, which makes a rigorous, assumption-explicit analysis valuable.

Key Objectives

- Identify what QKD does and does not provide, and where PQC is strictly required
- Design a hybrid key establishment architecture with an explicit key combiner
- Define which authentication mechanisms are appropriate and what their key management implies
- Compare the positions of national agencies on QKD and hybrid deployment
- Evaluate deployment scenarios and their threat models

Technical Stack:

- Python for simulation; QKD simulation libraries (SimulaQron, Qiskit QKD modules)
- Open source (liboqs, PQCclean) or internal PQC libraries
- Protocol modelling and formal analysis tooling

Expected Outcomes:

- A hybrid QKD/PQC architecture blueprint
- An explicit threat model and assumption list
- Simulated performance and key-rate data
- Recommendations for migration and for standards alignment

Team & Collaboration:

- Onboarded within the Advanced Research team, collaborating with cybersecurity, Lab attacks specialists.
- Potential for cross-functional integration with other secure communication initiatives.

Reference Papers

1. ETSI GS QKD 014 / QKD 015 — QKD key delivery API and control interface specifications
2. NSA, CNSA 2.0 and Quantum Key Distribution / Quantum Cryptography FAQ
3. ANSSI, BSI, NLNCSA, SE NCSA, *Position paper on quantum key distribution*
4. BSI, *Quantum-safe cryptography — fundamentals, current developments and recommendations*
5. NIST, SP 800-227, *Recommendations for Key Encapsulation Mechanisms* — hybrid key combiner guidance

[AST-AI-01] State-of-the-Art Exploration of AI World Models

Topics: **World Models, Foundation Models, ...**

Master Thesis Possible: **Yes**

World Models are emerging as a key building block of next-generation AI systems. By learning internal representations of the physical and digital world, these models can predict future states, simulate outcomes, support planning, and enable more capable autonomous agents. Major advances in multimodal AI, generative models, and embodied intelligence have accelerated research in this area, making World Models one of the most active fields in artificial intelligence today.

Details

The objective of this internship is to perform a comprehensive study of the state of the art in AI World Models and assess their potential for real-world deployment. The intern will investigate both academic and industrial developments, covering approaches such as latent world models, multimodal foundation models, video generation models, predictive architectures, embodied AI systems, and agent-based reasoning frameworks.

The internship will combine literature review, hands-on experimentation, technology benchmarking, and critical analysis. Selected open-source implementations will be explored and evaluated to better understand their capabilities, limitations, computational requirements, and applicability to autonomous systems and robotics.

The resulting work will contribute to identifying promising technologies and future development directions for intelligent autonomous platforms.

Key Objectives

- Conduct a comprehensive review of state-of-the-art World Model research.
- Analyze and compare major architectures from academia and industry.
- Evaluate applications in autonomous agents, robotics, planning, simulation, and decision making.
- Study the relationship between World Models, foundation models, and agentic AI systems.
- Benchmark selected open-source implementations and frameworks.
- Identify key technological challenges, limitations, and research opportunities.
- Produce recommendations and a technology roadmap for future developments.

Expected Outcomes:

- Comprehensive state-of-the-art survey and literature review.
- Comparative analysis of leading World Model architecture.
- Evaluation report covering strengths, weaknesses, and practical applicability of existing approaches.

-
- Experimental assessment of selected open-source implementations.
 - Identification of promising directions for future R&D activities.
 - Final technical report and presentation summarizing findings and recommendations.

Team & Collaboration:

- Onboarded within the Advanced Research team, collaborating with cybersecurity, Lab attacks specialists.
- Potential for cross-functional integration with other AI initiatives.

[AST-AI-02] State-of-the-Art Exploration of Audio Foundation Models for Detection and Situational Awareness

Topics: **Audio AI, Foundation Models, ...**

Master Thesis Possible: **Yes**

Recent advances in AI have enabled machines to understand and reason about audio with unprecedented accuracy. Modern audio foundation models can identify sounds, recognize speech, classify environments, detect anomalies, and extract meaningful events from both live and recorded audio streams. These capabilities are becoming increasingly important for robotics, security, monitoring, smart environments, and autonomous systems.

Details

The objective of this internship is to explore the state of the art in audio AI models and evaluate their ability to detect, classify, and interpret events from recorded and live audio streams.

The intern will investigate modern audio foundation models, acoustic event detection systems, speech and speaker recognition technologies, multimodal models, and emerging audio reasoning architectures. Particular emphasis will be placed on understanding how AI systems can detect meaningful events within complex real-world environments, such as alarms, machinery faults, vehicle activity, human interactions, impacts, breaking glass, animal sounds, emergency situations, and other operationally relevant signals.

The internship will combine literature review, technology benchmarking, and practical experimentation using open-source models and datasets. The goal is to identify the most promising technologies for deploying intelligent audio perception capabilities in future monitoring, robotics, and autonomous systems.

Key Objectives

- Conduct a comprehensive review of state-of-the-art audio AI and foundation models.
- Study acoustic event detection, sound classification, and audio scene understanding techniques.
- Evaluate speech recognition, speaker identification, and audio-language models.
- Benchmark selected open-source audio models on real-world detection tasks.
- Assess performance, latency, computational requirements, and scalability.
- Identify models suitable for live streaming and edge deployment scenarios.
- Investigate multimodal approaches combining audio with vision, language, or sensor data.
- Produce recommendations and a technology roadmap for future development.

Expected Outcomes:

- Comprehensive survey of state-of-the-art Audio AI technologies.
- Comparative evaluation of leading audio foundation models and detection architectures.

- Experimental benchmarks covering event detection, classification, and real-time processing capabilities.
- Assessment of suitability for live and recorded audio analysis applications.
- Identification of promising technologies for intelligent audio perception systems.
- Strategic recommendations and technology roadmap for future R&D activities.
- Final technical report and presentation summarizing findings and conclusions.

Team & Collaboration:

- Onboarded within the Advanced Research team, collaborating with cybersecurity, Lab attacks specialists.
- Potential for cross-functional integration with other AI initiatives.

[AST-AI-03] Synthetic World Simulation Platform

Topics: **Simulation, Synthetic Data, World Models, ...**

Master Thesis Possible: **Yes**

As AI and robotics systems become increasingly autonomous, the ability to train, test, and evaluate them in realistic virtual environments is becoming critical. Synthetic World Simulation Platforms enable the generation of scalable, controllable, and repeatable environments that can accelerate development while reducing the need for costly real-world data collection and testing.

Details

The objective of this internship is to explore, evaluate, and contribute to the development of a Synthetic World Simulation Platform capable of generating realistic virtual environments for AI training, validation, and experimentation.

The intern will investigate state-of-the-art simulation technologies, synthetic data generation pipelines, digital twin concepts, procedural world generation, and AI-driven environment creation. The project will assess how modern simulation platforms can support the training and evaluation of autonomous agents, robotics systems, perception models, world models, and decision-making algorithms.

The internship combines technology benchmarking, software development, and experimentation. The intern will evaluate existing simulation ecosystems, identify architectural requirements, and develop prototype capabilities that demonstrate how large-scale virtual worlds can be generated and leveraged for AI development.

Particular attention will be given to realism, scalability, reproducibility, sensor simulation, agent interaction, and integration with modern AI workflows.

Key Objectives

- Conduct a review of state-of-the-art simulation and synthetic world generation technologies.
- Analyze existing platforms used in AI, robotics, gaming, and digital twin applications.
- Evaluate procedural generation and AI-driven environment creation techniques.
- Investigate synthetic data generation for perception, planning, and autonomous systems.
- Design and prototype components of a scalable Synthetic World Simulation Platform.
- Assess realism, performance, scalability, and operational costs.
- Explore integration with world models, autonomous agents, and multimodal AI systems.
- Define recommendations and architecture guidelines for future development.

Expected Outcomes:

- Comprehensive assessment of modern simulation and synthetic-world technologies.
- Comparative evaluation of leading simulation platforms and frameworks.

-
- Prototype components demonstrating synthetic environment generation and agent interaction.
 - Analysis of realism, scalability, performance, and integration challenges.
 - Architectural recommendations for a future Synthetic World Simulation Platform.
 - Identification of opportunities for AI training, testing, validation, and synthetic data generation.
 - Final technical report and presentation summarizing findings, prototypes, and recommendations.

Team & Collaboration:

- Onboarded within the Advanced Research team, collaborating with cybersecurity, Lab attacks specialists.
- Potential for cross-functional integration with other AI initiatives.

[AST-AI-04] Context-Aware Robot using Local LLMs

Topics: **Local LLMs, Robotics, Edge AI, ...**

Master Thesis Possible: **Yes**

Large Language Models (LLMs) are enabling a new generation of robots capable of understanding natural language, reasoning about their environment, and interacting more effectively with humans. Running these models locally allows robots to operate with low latency, increased autonomy, and reduced dependence on cloud infrastructure, opening new possibilities for intelligent real-world applications.

Details

The objective of this internship is to explore, design, and evaluate a context-aware robotic architecture based on locally deployed Large Language Models.

The intern will investigate how robots can leverage local AI models to understand their environment, maintain contextual awareness, reason about situations, and interact naturally with users in dynamic environments. The project will assess the capabilities and limitations of modern open-source LLMs and multimodal models when deployed on robotic platforms.

Particular attention will be given to integrating language models with perception systems, memory modules, world models, planning components, and environmental sensors. The goal is to enable robots to build a richer understanding of their surroundings, remember relevant information, and make informed decisions based on context.

The internship combines research, prototyping, and experimentation, ultimately contributing to the development of next-generation autonomous robotic systems capable of natural and intelligent interactions.

Key Objectives

- Study state-of-the-art local LLM and multimodal AI architectures.
- Evaluate the feasibility of running advanced AI models directly on robotic hardware.
- Design a context-aware robotic architecture.
- Integrate local LLMs with perception, memory, planning, and decision-making components.
- Investigate mechanisms for maintaining situational and contextual awareness.
- Benchmark different open-source models and deployment strategies.
- Develop prototype capabilities demonstrating intelligent robot-human interaction.
- Explore the integration of world models and long-term memory systems.
- Define recommendations and architectural guidelines for future robotic platforms.

Expected Outcomes:

- Assessment of state-of-the-art local LLM technologies for robotics applications.
- Context-aware architecture for autonomous robotic systems.
- Prototype implementation integrating AI reasoning with robotic perception and decision-making.

- Benchmark results covering performance, latency, resource consumption, and scalability.
- Evaluation of contextual awareness and human-robot interaction capabilities.
- Recommendations for deploying local AI systems on robotic platforms.
- Final technical report and presentation summarizing architecture, experiments, findings, and future directions.

Team & Collaboration:

- Onboarded within the Advanced Research team, collaborating with cybersecurity, Lab attacks specialists.
- Potential for cross-functional integration with other AI initiatives.

[COE-AI-01] Agentic Web Monitoring for Audiovisual Piracy

Topics: **Agentic AI, web monitoring, multimodal analysis, web automation, audiovisual piracy**

Master Thesis Possible: **YES**

Live sports piracy remains a challenge for broadcasters and content owners, with infringing streams appearing and disappearing rapidly across media. Recent advances in AI agents, multimodal analysis, and automation create new opportunities to automate detection and evidence collection at scale. In this project, you will explore how agentic AI systems can be applied to monitoring of different sources and evaluate architectures that balance effectiveness, scalability, and operational cost. Based on the findings, you will contribute to the development of a production-ready monitoring platform.

Intern Mission

You will design, build, and evaluate components of an AI-powered monitoring platform for detecting audiovisual piracy. Working with experienced engineers, you will experiment with agentic architecture, develop monitoring and evidence-collection capabilities, and measure their effectiveness in realistic scenarios. Your work will contribute directly to the development of a scalable, production-ready system.

Key Objectives

- Design and evaluate agentic architectures for automated monitoring of sources.
- Develop scalable monitoring and evidence-collection capabilities.
- Measure effectiveness, scalability, and operational cost in realistic scenarios.
- Contribute components to a production-ready monitoring platform.

Expected Outcomes

- Evaluated agentic architectures for audiovisual-piracy monitoring.
- Monitoring and evidence-collection components tested in realistic scenarios.
- Systematic measurements to guide improvements in effectiveness, scalability, and operational cost.
- Contributions integrated toward a production-ready platform.

What You Will Learn

- How to design and evaluate agentic AI systems for real-world applications.
- How to build scalable data-collection and monitoring pipelines.
- How to measure and improve the performance of AI-driven systems through systematic experimentation.
- How production-grade systems are developed, tested, and operated.

[KS-MDR-01] MDR business operations & technology EPFL master's intern

Master Thesis Possible: **YES**

As an MDR Business Operations & Technology Masters Intern, you will complete a supervised industry project at the intersection of technology, management, and cybersecurity. Working with the MDR leadership team, you will apply analytical and management methods to a defined operational challenge and help translate business needs and technical information into clear, evidence-based recommendations.

Your Mission

You will support cross-functional work involving MDR operations, Security Engineering, Product & Platform, DevOps/IT, Delivery, Customer Success, Finance, and other business functions. Your assignment may include process analysis, performance measurement, business analytics, operating-model documentation, and evaluation of technology-enabled improvements. You will work with an assigned supervisor, present findings to stakeholders, and deliver practical outputs that can be adopted by the organization.

Your responsibilities will be:

1. Business analysis & stakeholder engagement

- Map stakeholder needs and document how business requests move into MDR operational or engineering work.
- Conduct structured interviews, workshops, and desk research under the guidance of the internship supervisor.
- Translate findings into concise problem statements, requirements, process maps, and decision-support materials.
- Support stakeholder meetings by preparing agendas, capturing decisions, and tracking agreed actions.
- Assess stakeholder pain points and identify improvement opportunities across business and technical interfaces.
- Prepare clear, data-supported updates and presentations for MDR leaders and project stakeholders.

2. Operations analysis, performance & continuous improvement

- Analyze an agreed intake-to-delivery process and establish a baseline using available operational data.
- Build simple dashboards or analytical models to assess lead time, backlog, readiness, quality, or other relevant measures.
- Apply quantitative methods to identify trends, bottlenecks, dependencies, and potential root causes.
- Compare improvement options using defined criteria such as customer impact, risk, effort, scalability, and expected value.

- Support the design or refinement of workflows, ownership models, readiness criteria, and operational controls.
- Document methods, assumptions, results, and recommendations so that the analysis can be reviewed and reused.

3. Technology strategy & implementation support

- Research relevant technologies, operating practices, and market developments affecting MDR service delivery.
- Assess the operational and business implications of proposed tools or process changes.
- Support pilots or proofs of concept by defining use cases, evaluation criteria, test plans, and feedback mechanisms.
- Create reusable project artifacts such as process documentation, analytical dashboards, stakeholder maps, implementation checklists, and presentation materials.
- Contribute to change and adoption planning, including communication, documentation, training inputs, and measurement of results.

4. AI-enabled SOC analysis

- Map the current workflow, user needs, data requirements, operational constraints, and human decision points for the selected use case.
- Define and analyze suitable measures, such as processing time, output quality, consistency, analyst effort, and user adoption.
- Present evidence-based findings, limitations, and recommendations while maintaining a human-in-the-loop approach.

You are

- Analytical and curious, with the ability to structure an ambiguous problem and test assumptions.
- A clear communicator who can connect technical findings with operational and business implications.
- Organized and dependable, able to manage a defined project, document progress, and meet agreed milestones.
- Collaborative and comfortable seeking feedback from specialists across different functions.
- Pragmatic, attentive to detail, and motivated to turn analysis into implementable recommendations.

You have

- Current enrollment in a relevant master's program combining engineering or science with management and technology, with eligibility to complete an industry-based master project or internship.
- A strong quantitative foundation, including statistics, data analysis, economics, operations, or comparable analytical coursework.
- Understanding of strategy, performance management, process improvement, and technology adoption in organizations.
- Ability to work with data using tools such as Excel, Power BI, Python, SQL, or equivalent analytical platforms.
- Strong written and verbal communication skills, including structured reports and presentations for technical and non-technical audiences.
- Interest in cybersecurity operations, AI-enabled workflows, or digital transformation; prior cybersecurity experience is an advantage but not required.
- Proficiency with common collaboration and planning tools, with disciplined documentation and project-management habits.
- Fluent in French and English; additional languages are a plus.

Why you'll love it here

If you are seeking a culture that supports growth, fosters success, and moves the industry forward, then Kudelski Security is where you need to be! As the premier provider of cybersecurity service and solutions, Kudelski Security has the most comprehensive ecosystem of security products and partners to deliver unparalleled services. Our rich and successful history with our clients is based on innovation and trust, serving clients of varying sizes and industries, including commercial, government and education. We have the proven expertise to plan, build, and run successful security programs across Data & Cloud security, Application Security, Microsoft 365 security, Data Privacy, Managed Security services, and IoT/OT security. Kudelski remains committed to advocating Diversity, Equality, and Inclusion within our organization and throughout the industry.

With Kudelski, you can expect

- Immersion in an incredible culture and the vibe of a fast moving and growing organization full of opportunity.
- Opportunity to work with talented peers.
- Creative problem-solving and the ability to tackle unique, complex projects.
- Competitive compensation with a benefits package that protects you and your loved ones and allows you to pursue career growth with tuition reimbursement.
- Generous time off for rest, relaxation, and hobbies.
- Colleagues from across the globe who are interested in helping clients protect their companies so they can focus on fulfilling their mission.

About Kudelski Security

Kudelski Security, a division of the Kudelski Group (SIX: KUD.S), is an innovative, independent Swiss provider of tailored cyber and media security solutions to enterprises and public sector institutions. Our team of security experts delivers end-to-end consulting, technology, managed services, and threat intelligence to help organizations build and run successful security programs. Our global reach and cyber solutions focus is reinforced by key international partnerships. Kudelski Security is headquartered in Phoenix, Arizona and Cheseaux-sur-Lausanne, Switzerland. For more information, please visit:

www.kudelskisecurity.com