CLUSTER SEQUOIA

Call for chairs 2026

Chairs submission

Ce travail a bénéficié d'une aide de l'Etat gérée par l'Agence Nationale de la Recherche au titre de France 2030 portant la référence ANR-23-IACL-0009.

CLUSTER SEQUOIA

JUNIOR CHAIRS

- **ACDC** : Anomaly Characterization, Detection and Classification on inspection robot for underwater structure, Christophe Viel (CNRS).
- **META** : MEsuring Trust in AI Systems, Romain Badouard (Inria)
- **MOSAIC** : Multi-robot Opinion-aware Socially Adaptive Intelligent Cohabitation, Esteban Restrepo (CNRS).
- **RAISE** : Reliable AI for The Next Generation of Safety- and MissionCritical Embedded Systems, Fernando Fernandes Dos Santos (Inria).
- **SAFE-AI Decisions** : Securing Acceptance of Fair and Explainable AI Decisions, Bertille Picard (ENSAI).
- **SCONE** : Self-supervised Removal of Spatially COrrelated NoisE, Sébastien Herbreteau (ENSAI).
- **SENTRY-AI** : SEcure aNd TRustworthY AI systems, Kallas Kassem (INSERM).
- **SYNTHETIC** : SYnthetic Network Traffic generation HEightened with Transfer learning for Intrusion detection systems and Cyber ranges, Pierre-François Gimenez (Inria).

EXPERIENCED CHAIRS

- **ACORD** : Algorithms for Combinatorial Optimization and Representation learning with Differentiable solvers, Romain Tavenard (Université Rennes 2).
- **CLU** : On-Chip Uncertainty-Aware Continual Learning, Jean-Philippe Diguet (CNRS).
- **DOLPHIN-DP** : Decentralized and OnLine Privacy-preserving macHIne learNing through Differential Privacy, Romaric Gaudel (Université de Rennes).
- **GRASP** : GRAph Signal Processing for Interpreting and Steering Multimodal LLMs, Nicolas Farrugia (IMT Atlantique).
- **RECONFIAL** : AI and strategic reconfigurations for an international-local continuum of governance, Sandrine Turgis (Université de Rennes).
- **VISTORIA** : Visual and Interactive Storytelling for AI Explainability, Luis Galarraga Del Prado (Inria).

INTERNATIONAL CHAIRS

- **AI-STAR** : Artificial Intelligence for Structured Threat Modeling with Attack Trees, Sophie Pinchinat (Université de Rennes).
- **AI4MicroSecurity** : AI-driven Security Against Microarchitectural Vulnerabilities, Guy Gogniat (Université de Bretagne Sud).
- **FEDULCAD** : FEDerated and frUgal Learning for Complex Attack Detection, David Espes (Université de Bretagne Occidentale).
- **PYTHAI** : PrivacY-preserving and TrustwortHy explainable AI, Tristan Allard (Université de Rennes).

CHAIR PROPOSAL CHAIR Junior / ACDC

CALL FOR CHAIRS 2026

Acronym	ACDC
Chair Title	**Anomaly Characterization, Detection and Classification on inspection robot for underwater structure**
SequoIA research pilar	☐ Pilar 1: Core AI ☒ Pillar 2: AI, cybersecurity and defense ☒ Pillar 3: AI for environment, ocean
Chair type	☒ Junior ☐ Experienced
Keywords	AI, Inspection, anomaly detection
Expected beginning / duration	September 2026, 42 months
Leading institution	Centre National de la Recherche Scientifique – Lab-STICC (UMR6285)
☒ **Please confirm you have informed your Head of Department of your submission**	
Chair Holder	**Last Name, First Name, Position** VIEL, Christophe, CR CNRS – Lab-STICC – ENSTA **e-mail** [redacted] **Phone** [redacted]
Co-Chair Holder (if any)	**Last Name, First Name, Position** BERGANTIN, Lucia, IMT Atlantique, Lab-STICC, Associate Professor **e-mail** [redacted] **Phone** [redacted]

Abstract of the chair project (Non-Confidential – 4000 characters maximum, including spaces). This section will be shared with institutional and industrial partners of the Cluster

Underwater marine structures like ports, offshore wind turbines, and submerged channels require regular inspections to detect sabotage, espionage, and early signs of degradation. Factors such as depth, cold temperatures, polluted waters, rough seas, and prolonged inspection times make human diving difficult and dangerous. Moreover, deploying sensors across such large infrastructures is extremely costly and requires continuous maintenance. Underwater robots offer a more viable solution for inspecting such infrastructures replacing lengthy, on-site inspections by humans with autonomous anomaly detection and reporting by a robot, with the aim of assisting human operators in the generation of mission reports. This requires intelligent perception to detect and localize anomalies such as unidentified objects (e.g., mines), cracks, corrosion, deposits, or

other type of degradation. Due to turbidity and low light, acoustic sensors are more often employed, reserving the use of conventional optical cameras for clearer or emerged sections.

The objective of this Chair is to advance the state-of-the-art in autonomous underwater inspection by focusing on robust in-the-wild detection and localization of structural anomalie, despite adverse conditions and evolving target appearances in underwater environments. These new capacities will be embedded to an experimental underwater robot to perform routine or on-demand inspections over long distances or large areas, in submerged environments that are largely or completely inaccessible to humans. The axis presented here focuses on the detection of anomalies using Artificial Intelligence (AI). Other aspects of the project, such as robot control, location tracking, etc., will be addressed through other funding applications.

From an AI perspective, the Chair will leverage neural networks as universal function approximators for the following tasks: representing the shape and structure of anomalies, and the perception (detection & classification) of anomalies. The diversity in scale, shape and category of anomalies and the structures that can be inspected by a robot (e.g. dams, offshore wind turbines, canals, tunnels, submerged caves, mines, or boat hulls), calls for a data-driven approach for object representation and perception. Embedding such models in underwater inspection robots (constrained by limited computing power and the use of acoustic imaging) makes off-the-shelf, state-of-the-art solutions either incompatible or suboptimal. Likewise, the absence/scarcity of sonar datasets containing anomalies inhibits model training. The Chair will address these challenges via transfer learning and domain adaptation techniques to leverage existing optical datasets, combined with data augmentation techniques specifically tailored to underwater images (e.g., synthetic and real-image mixing and filter-based). We plan to build on [11] and our previous study [1] to develop our new anomaly detection pipeline. Physics-informed AI strategies should also be explored, to take into consideration spatial and physical relationships.

Anomaly detection on sonar images will be the subject of a PhD thesis, co-supervised between the team ROBEX (specialized in underwater robotics and acoustic cameras) and RAMBO (specialized in AI) of the Lab-STICC laboratory. Experiments will first be conducted in ENSTA's testing pool, and later under more realistic conditions in a port, a lake and in an underwater pipe. Tests in pipes will be possible thanks to the collaboration with the industrial partner Sub-C-Marine, which will give us access to the European testing center C.E.E.R.I.C. (Centre d'Essai Europeen des Réseaux Industriels & Conduite). Sub-C-Marine will also provide data to validate the model, and is interested in implementing our methods on their robots for pipe inspection.

CLUSTER SEQUOIA

CALL FOR CHAIRS 2026

CHAIR PROPOSAL: Measuring Trust in AI Systems/ META

Acronym	META
Chair Title	MEsuring Trust in AI Systems
SequoIA research pilar	☒ Pilar 1: Core AI ☐ Pillar 2: AI, cybersecurity and defense ☐ Pillar 3: AI for environment, ocean
Chair type	☒ Junior ☐ Experienced
Keywords	Trust metrics ; AI standards ; AI Governance ; AI uses ; Democratic experiments
Expected beginning / duration	3 years from October 1st, 2026
Leading institution	Inria
☒ **Please confirm you have informed your Head of Department of your submission**	

Chair Holder	**Last Name, First Name, Position**	
	Badouard, Romain, Inria Researcher	
	e-mail	**Phone**
	[redacted]	[redacted]
Co-Chair Holder (if any)	**Last Name, First Name, Position**	
	e-mail	**Phone**

Abstract of the chair project (Non-Confidential – 4000 characters maximum, including spaces). This section will be shared with institutional and industrial partners of the Cluster

Measuring trust is both central to AI regulation practices and to the dynamics of AI adoption. Yet, the available metrics, indicators, and scales suffer from well-documented limitations. Drawing on social science approaches rooted in Science and Technology Studies (STS) and the sociology of quantification, the META Chair aims both to study how "trust in AI" is being transformed into a measurable and governable object through assessment frameworks, and to explore alternative ways of constructing trust indicators inspired by participatory methods and human-computer interaction design. By collaborating with standardisation bodies and European regulators, the project seeks to join European scientific networks dedicated to AI governance, while developing international partnerships around issues of trust and over-trust in AI systems—topics that lie at the core of current debates on AI Safety. Research conducted within the Chair will also provide cluster partners with reviews of existing trust metrics in AI and propose alternative indicators, tested in experimental settings, that can be implemented across various professional contexts.

CLUSTER SEQUOIA

CALL FOR CHAIRS 2026

CHAIR PROPOSAL Multi-robot Opinion-aware Socially Adaptive Intelligent Cohabitation / MOSAIC

Acronym	MOSAIC
Chair Title	Multi-robot Opinion-aware Socially Adaptive Intelligent Cohabitation
SequoIA research pilar	☒ Pilar 1: Core AI ☒ Pillar 2: AI, cybersecurity and defense ☐ Pillar 3: AI for environment, ocean
Chair type	☒ Junior ☐ Experienced
Keywords	Human-AI and Human-Robot interaction; Multi-Agent Systems; Autonomous robotics and interactions
Expected beginning / duration	September 2025 / 48 months
Leading institution	CNRS
☒ **Please confirm you have informed your Head of Department of your submission**	

Chair Holder	**Last Name, First Name, Position**	
	RESTREPO, Esteban, CR CNRS at IRISA	
	e-mail	**Phone**
	[redacted]	[redacted]
Co-Chair Holder (if any)	**Last Name, First Name, Position**	
	e-mail	**Phone**

Abstract of the chair project (Non-Confidential – 4000 characters maximum, including spaces). This section will be shared with institutional and industrial partners of the Cluster

Multi-Robot Systems (MRS) are highly effective for executing collaborative tasks, but their transition to complex, unstructured, and human-shared environments across various physical domains remains a significant challenge. While Artificial Intelligence (AI) has driven major progress in single-agent autonomy, the deployment of multi-agent teams dynamically interacting with human actors is currently hindered by limitations in distributed decision-making, intention understanding, and adaptability.

The MOSAIC (Multi-robot Opinion-aware Socially Adaptive Intelligent Cohabitation) project seeks to overcome these barriers by achieving "Socially-Aware MRS Cohabitation." MOSAIC fundamentally departs from the traditional robotics paradigm that treats human-shared spaces merely as dynamic obstacle courses. Instead, it models complex social dynamics as continuously evolving "opinions." To implement this, the project develops a unified, socially-aware control framework that seamlessly integrates Multi-Agent Reinforcement Learning (MARL) with continuous nonlinear opinion dynamics.

By representing the robotic team as distributed agents interacting over a complex graph, the framework will optimize opinion-dynamics-based models in real time utilizing MARL-based policies. This novel synthesis combines a formal mathematical structure for modeling passive human social influence with the real-time adaptability required to continuously optimize the MRS's collective behavior. Consequently, the robotic team moves beyond implicit intention estimation to actively use emergent belief formation, enabling it to safely cohabitate with human actors while fulfilling global mission specifications.

The project is structured around three sequential core activities. The first activity focuses on the formal mathematical modeling and analysis of human and multi-robot interactions using continuous nonlinear opinion dynamics integarted with MARL. Building upon this theoretical foundation, the second activity centers on the construction of a MARL-driven optimization architecture, leveraging decentralized training with decentralized execution (DTDE), Liquid-Graph Time-constant (LGTC) networks for scalability, and safe exploration strategies to adapt multi-robot interaction graphs and parameters in real time. Finally, the third activity is dedicated to physical validation in a demonstrator. During this stage, a robotic team will perform a mock-up mission, such as environmental monitoring or resource distribution, to prove its ability to dynamically adapt to the physical presence of humans in a shared space.

Ultimately, this model-informed learning approach provides formal stability and safety guarantees for human-MRS cohabitation. By bridging the gap between theoretical intention modeling and real-time distributed control, MOSAIC aligns directly with the Pillars (1) Core AI and (2) Cybersecurity/Defense of the SequoIA roadmap. It establishes a robust foundation for the safe integration of autonomous systems into human-centric societies, enhancing the autonomy and resilience of robots in critical real-world applications such as disaster response, precision agriculture, and construction.

Cluster SequoIA
Call for chairs
2025

CHAIR PROPOSAL

FERNANDO FERNANDES DOS SANTOS/RAISE

Acronym	Fernando Fernandes dos Santos/RAISE	
Chair Title	Reliable AI for The Next Generation of Safety- and Mission-Critical Embedded Systems	
SequoIA research pilar	☒ Pilar 1: Core AI ☒ Pillar 2: AI for cybersecurity and defense ☐ Pillar 3: AI for environment and ocean	
Chair type	☒ Junior ☐ Experienced	
Keywords		
Expected beginning / duration	48 months	
Leading institution		
☐ **Please confirm you have informed your Head of Department of your submission**		
Chair Holder	**Last Name, First Name, Position**	
	FERNANDES DOS SANTOS, Fernando, ISFP	
	e-mail	**Phone**
	[redacted]	[redacted]
Co-Chair Holder (if any)	**Last Name, First Name, Position**	
	e-mail	**Phone**

Cluster SequoIA
Call for chairs
2025

CHAIR PROPOSAL

FERNANDO FERNANDES DOS SANTOS/RAISE

Summary of the chair project (Non-Confidential – 4000 characters maximum, including spaces)

In safety- and mission-critical embedded systems, performance, power, and reliability are equally important design constraints. This is especially true for large AI models deployed in critical embedded systems, such as autonomous driving and aerospace, where inference must operate under strict energy budgets and harsh conditions, including radiation, temperature extremes, and long operational lifetimes.Commercial off-the-shelf AI hardware accelerators offer high computational efficiency but limited reliability protections. In contrast, hardened platforms offer greater resilience at the cost of performance and energy efficiency, making both options insufficient for modern, large-scale transformer-based inference workloads. At the software level, existing reliability assessment and fault-tolerance methods were largely developed for smaller models and do not scale well to large AI models, nor do they generally provide analytical frameworks that relate low-level hardware faults to incorrect inferences.

These limitations motivate a cross-layer approach in which hardware faults are traced through model operations and hardware structures to identify the subset that truly affects inference. *Reliable AI for the Next Generation of Safety- and Mission-Critical Embedded Systems* (RAISE) project is then grounded in the hypothesis that not all faults matter for modern large AI models, and that this information can be exploited for more efficient reliability assessment and fault mitigation. Building on this view, RAISE will develop scalable methods for reliability assessment and fault mitigation based on fault traceability and sensitivity-guided selective protection, bridging AI models and hardware accelerators for dependable AI deployment in critical embedded applications. Compared with prior work, RAISE focuses on large AI models and cross-layer methods that connect hardware faults to incorrect inferences. For experimental validation, RAISE will combine microarchitectural fault simulation with realistic fault-injection campaigns, including radiation and temperature experiments, to quantify failure rates and assess mitigation methods under representative operating conditions. RAISE aims to make reliability assessment tractable for large AI models by combining AI-level sensitivity metrics with microarchitectural fault analysis. RAISE will target failure-rate objectives consistent with functional-safety requirements for terrestrial applications and with acceptable failure rates for space applications, defined jointly with industrial partners.

The core innovation of RAISE is to make AI reliability predictive rather than exhaustive, by linking hardware fault locations to AI model failures and using that link to drive selective cross-layer protection.

CALL FOR CHAIRS 2026

CHAIR PROPOSAL Securing Acceptance of Fair and Explainable AI Decisions / **SAFE-AI Decisions**

Acronym	SAFE-AI Decisions
Chair Title	Securing Acceptance of Fair and Explainable AI Decisions
SequoIA research pilar	☒ Pilar 1: Core AI ☐ Pillar 2: AI, cybersecurity and defense ☐ Pillar 3: AI for environment, ocean
Chair type	☒ Junior ☐ Experienced
Keywords	Human-AI interaction, decision making, algorithmic explainability
Expected beginning / duration	September 2026 / 4 years
Leading institution	Ensai
☒ **Please confirm you have informed your Head of Department of your submission**	

Chair Holder	**Last Name, First Name, Position**	
	Picard, Bertille, Assistant Professor (tenure-track position)	
	e-mail	**Phone**
	[redacted]	[redacted]
Co-Chair Holder (if any)	**Last Name, First Name, Position**	
	e-mail	**Phone**

Abstract of the chair project (Non-Confidential – 4000 characters maximum, including spaces). This section will be shared with institutional and industrial partners of the Cluster

This project aims to better understand the cognitive and behavioral mechanisms that determine individuals' acceptance or aversion to algorithmic recommendations, in order to identify the conditions under which Artificial Intelligence (AI)-assisted decisions are accepted. It focuses in particular on how decision-makers trade off predictive performance, explainability and fairness, concepts that are widely hypothesized to foster trust in algorithms in the machine learning (ML) literature. However, while numerous theoretical studies have proposed methods to make models more transparent, their actual impact on human behavior remains debated. Moreover, while the ML literature has extensively studied these concepts from a technical perspective, their alignment with decision-makers' expectations remains poorly understood, which may hinder the adoption of such tools in industrial and public decision-making contexts.

CALL FOR CHAIRS 2026

CHAIR PROPOSAL Securing Acceptance of Fair and Explainable AI Decisions / **SAFE-AI Decisions**

The project proposes an approach to the acceptability of automated decision-making systems that explicitly takes user preferences into account, particularly with respect to explainability and fairness. This issue is all the more central given that European regulatory frameworks, such as GDPR and the AI Act, emphasize the importance of transparency in algorithmic systems affecting individuals.

The project adopts an interdisciplinary approach at the intersection of ML, economics and psychology. It relies on three complementary components: (1) the development of a theoretical framework to model how individuals evaluate AI-assisted decision systems, (2) the collection of survey data to empirically characterize decision-makers' preferences regarding explainability and fairness across different contexts, and (3) laboratory experiments testing whether decision systems aligned with these preferences generate higher levels of trust and acceptance than systems that do not.

Applications are wide-ranging, particularly in domains where algorithmic decisions carry significant consequences for individuals, such as public policy allocation, medical treatment decisions, or insurance pricing for instance.

By combining theoretical modelling with empirical and experimental approaches, the SAFE-AI Decisions project will contribute to a better understanding of the conditions under which AI-assisted decisions are accepted and provide insights for designing algorithmic systems that better align user expectations.

CLUSTER SEQUOIA

Call for chairs 2026

CHAIR PROPOSAL

SCONE

Acronym	SCONE
Chair Title	Self-supervised Removal of **S**patially **CO**rrelated **N**ois**E**
SequoIA research pilar	☒ Pilar 1: Core AI ☒ Pillar 2: AI for cybersecurity and defense ☐ Pillar 3: AI for environment and ocean
Chair type	☒ Junior ☐ Experienced
Keywords	image denoising, correlated noise, weakly-supervised learning, transfer learning
Expected beginning / duration	September 2026 / 4 years
Leading institution	ENSAI / CREST UMR 9194
☒ **Please confirm you have informed your Head of Department of your submission**	
Chair Holder	**Last Name, First Name, Position** HERBRETEAU Sébastien, Assistant Professor **e-mail** [redacted] **Phone** [redacted]
Co-Chair Holder (if any)	**Last Name, First Name, Position** **e-mail** **Phone**

Summary of the chair project

We propose to address an underexplored problem in image denoising: self-supervised learning of neural networks for the removal of spatially correlated noise. Indeed, the image denoising literature is vast and offers a wide range of methods—both supervised and self-supervised—for handling spatially independent noise. In particular, denoising images corrupted by additive white Gaussian noise is arguably one of the most extensively studied inverse problems to date, alongside Poisson–Gaussian noise, which has driven significant advances in fields such as fluorescence microscopy and satellite imaging. In recent years, the field has evolved considerably with the emergence of self-supervised approaches, which eliminate or reduce the need for clean training data. This is especially valuable in settings where acquiring ground-truth images is either time-consuming or entirely infeasible. These methods typically rely on carefully designed loss functions or architectural constraints embedded directly into the neural networks.

However, spatially correlated noise remains much less studied, despite being highly relevant in practice. Such noise naturally arises in many real-world imaging systems due to physical acquisition processes, sensor imperfections, or reconstruction pipelines. Ignoring these correlations can lead to suboptimal denoising performance, loss of fine structures, or the introduction of artifacts, making it crucial to explicitly account for them. We aim to extend self-supervised denoising techniques to this more challenging setting. Existing approaches mainly fall into two categories, each with notable limitations. The first attempts to break spatial correlations—for instance, through downsampling—before applying methods designed for independent noise, inevitably leading to a loss of information. The second, which appears more promising, combines unbiased estimators coupled with constraints within the network, though it often relies on restrictive assumptions such as a known or constant noise level.

We propose to investigate a third, hybrid direction: weakly supervised learning, based on adapting a denoiser trained on synthetic data and transferring it efficiently to real data. This adaptation would rely on lightweight strategies inspired by plug-and-play methods or parameter-efficient fine-tuning. The project will aim to formalize these ideas, design suitable architectures, and systematically evaluate their performance across a variety of datasets. Ultimately, our goal is to demonstrate that effective alternatives to fully supervised learning can be developed for spatially correlated noise, opening the door to more realistic and widely applicable denoising methods.

CALL FOR CHAIRS 2026

Acronym	SENTRY-AI
Chair Title	SENTRY-AI: Secure and Trustworthy AI Systems – From Federated Learning to Foundation Models
SequoIA research pilar	☒ Pilar 1: Core AI ☒ Pillar 2: AI, cybersecurity and defense ☐ Pillar 3: AI for environment, ocean
Chair type	☒ Junior ☒ Experienced
Keywords	Trustworthy AI, AI Security, Federated Learning, Foundation Models, Backdoor Attacks, Cybersecurity.
Expected beginning / duration	Expected beginning: September 2026 Duration: 48 months (4 years)
Leading institution	INSERM – LaTIM (Laboratory of Medical Information Processing)
☒ **Please confirm you have informed your Head of Department of your submission**	
Chair Holder	**Last Name, First Name, Position**: Kassem, Kallas, Senior Scientist **e-mail**: [redacted] **Phone**: [redacted]
Co-Chair Holder (if any)	**Last Name, First Name, Position**: **e-mail**: **Phone**:

Abstract of the chair project (Non-Confidential – 4000 characters maximum, including spaces). This section will be shared with institutional and industrial partners of the Cluster

The rapid deployment of artificial intelligence (AI) systems in critical domains such as healthcare and cybersecurity has introduced unprecedented challenges related to trust, robustness, and security. In particular, federated learning and foundation models are reshaping modern AI pipelines, enabling large-scale collaborative learning while simultaneously expanding the attack surface through complex, distributed, and often opaque mechanisms. This chair is a vital

CLUSTER SEQUOIA

CALL FOR CHAIRS 2026

technological enabler for achieving mandatory legal/GDPR compliance, establishing clinical trust, and securing social acceptability for AI in high-stakes healthcare deployments.
This project, SENTRY-AI, aims to establish a unified scientific framework for the security and trustworthiness of distributed and foundation AI systems under realistic adversarial conditions. The project addresses fundamental challenges at the intersection of core AI and cybersecurity by investigating how learning systems behave, adapt, and fail when exposed to adaptive, coordinated, and stealthy adversaries.

The research is structured around three complementary axes. First, we study the theoretical and algorithmic foundations of adversarially robust learning in distributed and non-IID environments, with a focus on understanding the limits of robustness in federated and large-scale models. Second, we design novel attack and defense mechanisms targeting federated learning and foundation-model pipelines, including backdoor attacks, supply-chain vulnerabilities, and parameter-efficient adversarial manipulations, with a particular emphasis on healthcare applications. Third, we develop a unified benchmarking and validation framework to systematically evaluate security risks and defense mechanisms, ensuring reproducibility, comparability, and real-world relevance.

By combining theoretical insights, system-level innovations, and application-driven validation, SENTRY-AI will contribute to the development of trustworthy AI systems that can be deployed safely in high-stakes environments. The project will reinforce the scientific objectives of Cluster SequoIA by advancing research at the interface of core AI and cybersecurity, while fostering interdisciplinary collaborations and promoting open, reproducible science.

CALL FOR CHAIRS 2026

CHAIR PROPOSAL CHAIR NAME / ACRONYM

SYNTHETIC

Acronym	SYNTHETIC
Chair Title	SYnthetic Network Traffic generation HEightened with Transfer learning for Intrusion detection systems and Cyber ranges
SequoIA research pillar	☐ Pilar 1: Core AI ☒ Pillar 2: AI, cybersecurity and defense ☐ Pillar 3: AI for environment, ocean
Chair type	☒ Junior ☐ Experienced
Keywords	AI for Cybersecurity, AI-based generation, Transfer learning, Cyber ranges, Intrusion detection system
Expected beginning / duration	Starting October 2026 for 36 months
Leading institution	Inria
☒ **Please confirm you have informed your Head of Department of your submission**	
Chair Holder	**Last Name, First Name, Position** GIMENEZ Pierre-François, Inria Research Scientist (ISFP) **e-mail** [redacted] **Phone** [redacted]
Co-Chair Holder (if any)	**Last Name, First Name, Position** **e-mail** **Phone**

Abstract of the chair project (Non-Confidential – 4000 characters maximum, including spaces). This section will be shared with institutional and industrial partners of the Cluster

SYNTHETIC aims at improving synthetic network traffic generation through transfer learning. Network traffic is essential for evaluating cybersecurity solutions, such as the false-positive rates of intrusion detection systems (IDS), and for increasing the realism of honeypots and cyber ranges. Current methods for collecting data, either from real systems or experimental platforms, have issues with privacy, realism, and obsolescence. Recently, methods relying on artificial intelligence and machine learning have been proposed to generate synthetic network traffic. However, these methods have a major limitation: they cannot be parametrized to generate traffic for any network description. In SYNTHETIC, we propose to lift that limitation through transfer learning with two main contributions: unsupervised transfer learning of benign network traffic generation and style transfer of malicious traffic. We propose then to exploit cross-domain generation to enhance IDS and cyber ranges. Currently, there is no way to evaluate the robustness of IDS against dataset distribution shift because the datasets are stationary. With a dataset generation method that can be parametrized by a network description, we plan to generate datasets with gradual or

sudden distribution shifts, such as the addition or deletion of a server or a technology, the modification of network topology, etc. Such datasets could foster research into improving the stability of IDS performance over time, rather than assuming that a model should be retrained periodically. Cyber ranges and honeypots encompass various network topologies depending on the scenario, so a generator that can adapt to any network description could generate background traffic in any circumstance. Background traffic is essential in honeypots so attackers cannot easily detect that they have infiltrated a fake network without any actual users or data. It is also essential in cyber ranges, making the exercise more realistic for the defense team (the “blue” team) and, therefore, the learned skills more transferable to actual cyber crises.

Cluster SequoIA
Call for chairs
2025

CHAIR PROPOSAL

ACORD

Acronym	ACORD	
Chair Title	Algorithms for Combinatorial Optimization and Representation learning with Differentiable solvers	
SequoIA research pilar	☒ Pilar 1: Core AI ☐ Pillar 2: AI for cybersecurity and defense ☐ Pillar 3: AI for environment and ocean	
Chair type	☐ Junior ☒ Experienced	
Keywords	time series ; foundation models ; combinatorial optimization ; optimal transport	
Expected beginning / duration	October 2026 / 48 months	
Leading institution	Univ. Rennes 2	
☒ **Please confirm you have informed your Head of Department of your submission**		
Chair Holder	**Last Name, First Name, Position**	
	TAVENARD Romain, Professor	
	e-mail	**Phone**
	[redacted]	

Cluster SequoIA
Call for chairs
2025

CHAIR PROPOSAL

ACORD

Summary of the chair project (Non-Confidential 4000 characters maximum, including spaces)

Recent progress in machine learning has been driven by the integration of increasingly complex algorithmic components into end-to-end differentiable systems. However, many fundamental operations underlying modern data analysis such as sequence alignment, graph optimization, and optimal transport still rely on discrete algorithms that are difficult to integrate into gradient-based learning pipelines. We aim to develop differentiable solvers for two main reasons: (i) they can reduce computational cost at the cost of approximations; (ii) they can serve as building blocks in deep learning pipelines. The project applies this agenda in two domains: time series forecasting (WP2) and scalable optimal transport (WP3).

First, the project will develop differentiable solvers for combinatorial optimization problems, with a particular emphasis on shortest-path computations and optimal transport. These methods will rely on continuous relaxations and gradient-based optimization to produce high-quality approximate solutions while preserving differentiability, enabling their integration within neural architectures.

Building on these foundations, the project will target two main areas: time series forecasting and scalable optimal transport. For the former, it will investigate new modeling approaches for time series forecasting, with a focus on irregularly sampled signals and continuous-time representations, for which alignment-based losses deriving from the differentiable solvers will be employed. In parallel, it will contribute to the development of reproducible benchmarks for time series foundation models, helping clarify their current capabilities and limitations.

Finally, the project will explore a new paradigm for optimal transport based on learning end-to-end transport solvers. Building on sliced optimal transport methods, we will investigate architectures capable of learning how to construct efficient transport plans across heterogeneous data distributions.

By combining advances in differentiable optimization, time series modeling, and optimal transport, the project aims to establish new algorithmic tools for machine learning on structured data. Beyond methodological contributions, these developments have the potential to impact a wide range of applications, including generative modeling, distributional learning, and large-scale analysis of temporal data.

CALL FOR CHAIRS 2026

CHAIR PROPOSAL On-Chip Uncertainty Aware Continual Learning / CLUE

Acronym	CLU
Chair Title	On-Chip Uncertainty-Aware Continual Learning
SequoIA research pilar	☒ Pilar 1: Core AI ☐ Pillar 2: AI, cybersecurity and defense ☐ Pillar 3: AI for environment, ocean
Chair type	☐ Junior ☒ Experienced
Keywords	In-Memory computing, Uncertainty estimation, Adap
Expected beginning / duration	Oct. 2026 / 4 years
Leading institution	CNRS

☒ **Please confirm you have informed your Head of Department of your submission**

Chair Holder	**Last Name, First Name, Position**	
	DIGUET, Jean-Philippe, DR CNRS, Lab-STICC	
	e-mail	**Phone**
	[redacted]	[redacted]

Co-Chair Holder (if any)	**Last Name, First Name, Position**	
	e-mail	**Phone**

***Abstract**: The Chair CLUE will investigate new research directions to make possible the emergence of tiny devices that can self-monitor the quality or uncertainty of their results and use this information to keep on learning with their own resources without connection to the cloud. It will explore new approaches of self-adaptive micro adaptive ensemble of models, new concepts of in-memory computing for the inference of binary or ternary neural networks provide and will provide new features to facilitate the estimation of uncertainty and the control of the re-training process. The CLUE team will be multi-disciplinary and international involving experts from France, Australia and Japan in memory and processor architectures, embedded smart systems, machine learning including uncertainty estimation, ternary logic dealing uncertain states, advanced non-volatile memory technologies and autonomous systems.*

CLUSTER SEQUOIA

CALL FOR CHAIRS 2026

CHAIR PROPOSAL Romaric GAUDEL / DOLPHIN-DP

Acronym	DOLPHIN-DP
Chair Title	Decentralized and OnLine Privacy-preserving macHIne learNing through Differential Privacy
SequoIA research pilar	☒ Pilar 1: Core AI ☒ Pillar 2: AI, cybersecurity and defense ☐ Pillar 3: AI for environment, ocean
Chair type	☐ Junior ☒ Experienced
Keywords	Differential Privacy, Decentralized Learning, Multi-Armed Bandits
Expected beginning / duration	January 2027 / 4 years
Leading institution	IRISA / Univ Rennes
☒ **Please confirm you have informed your Head of Department of your submission**	

Chair Holder	**Last Name, First Name, Position**	
	GAUDEL, Romaric, Maître de Conférence	
	e-mail	**Phone**
	[redacted]	[redacted]
Co-Chair Holder (if any)	**Last Name, First Name, Position**	
	e-mail	**Phone**

Abstract of the chair project (Non-Confidential – 4000 characters maximum, including spaces). This section will be shared with institutional and industrial partners of the Cluster

Artificial Intelligence systems are increasingly deployed in decentralized, adaptive, and adversarial environments such as critical infrastructures, distributed sensor networks, healthcare consortia, and defense systems. In these contexts, trust is not an abstract virtue but a measurable property: it depends on whether privacy guarantees are mathematically sound, whether attacker models are realistic, and whether security mechanisms are optimized rather than heuristic. Privacy and cybersecurity are therefore inseparable. A privacy guarantee that ignores communication topology may be misleading; a security mechanism that ignores statistical inference may be fragile.

While differential privacy (DP) provides a gold standard for privacy guarantees, its current applicability remains limited by the specific scenarios for which it was developed. For instance, in decentralized learning (DL), existing DP works are restricted to a narrow set of attacker models,

while protection strategies remain ad-hoc for each setting; although DL framework opens new avenues for attacks, it lacks a unified defense mechanism. Similarly, approaches considering sequential learning are rarely equipped with rigorous DP protection. Finally, privacy analysis is typically bounded by worst-case assumptions that ignore correlations in the data exposed, leading to conservative choices that degrade learning utility. In this project, we will make advances on all these aspects. While we target fundamental advances, they will enable true applicability of DP especially in cybersecurity and defense contexts.

CHAIR PROPOSAL **GRA**ph **S**ignal **P**rocessing for Interpreting and Steering Multimodal LLMs. **/ GRASP**

CALL FOR CHAIRS 2026

Acronym	GRASP
Chair Title	**GRA**ph **S**ignal **P**rocessing for Interpreting and Steering Multimodal LLMs.
SequoIA research pilar	X Pillar 1: Core AI ☐ Pillar 2: AI, cybersecurity and defense ☐ Pillar 3: AI for environment, ocean
Chair type	☐ Junior X Experienced
Keywords	Interpretability, graph signal processing, multimodal, LLM, neuroimaging, ecoacoustics
Expected beginning / duration	2026 / 4 years
Leading institution	IMT Atlantique
X **Please confirm you have informed your Head of Department of your submission**	

Chair Holder	**Last Name, First Name, Position**	
	Farrugia, Nicolas , Full Professor	
	e-mail	**Phone**
	[redacted]	[redacted]
Co-Chair Holder (if any)	**Last Name, First Name, Position**	
	e-mail	**Phone**

Abstract of the chair project (Non-Confidential – 4000 characters maximum, including spaces). This section will be shared with institutional and industrial partners of the Cluster

As Artificial Intelligence (AI) transitions into high-stakes domains—such as medical applications or environmental monitoring—the lack of transparency in its decision-making remains a critical barrier to adoption. While the field of mechanistic interpretability has successfully identified "atomized" features using Sparse Autoencoders (SAEs), moving from individual neurons to a global understanding of the computational circuits governing complex reasoning remains an open challenge. This is particularly acute in Multimodal Large Language Models (MLLMs), which frequently suffer from "perceptual decoupling" (modality bias), where the model prioritizes internal linguistic priors over external sensory data such as

CLUSTER SEQUOIA

CALL FOR CHAIRS 2026

EEG, fMRI, or acoustic signals. The GRASP project proposes a transformative framework for AI interpretability by introducing Graph Signal Processing (GSP) as a formal proxy for a model's internal computational state. By treating neural architectures as communication graphs and activations as signals, GRASP moves beyond local feature discovery to characterize the spectral dynamics of reasoning. We leverage a fundamental analogy from cognitive neuroscience—Structure-Function (SF) coupling—to quantify and steer the alignment between a model's physical topology and its functional output. To achieve this, we will first abstract modern architectures like Transformers and State-Space Models into formal graphs. This enables the definition of a Graph Fourier Transform of neural activations, allowing us to decompose complex reasoning into interpretable spectral components and define a rigorous metric for representational quality. Building on this foundation, we will apply GSP metrics to identify perceptual decoupling in MLLMs and introduce steering methods based on graph signal filtering. Such techniques apply spectral graph filters to activations to suppress biased noise and recouple the model to its sensory inputs without the need for costly retraining. The framework will be validated through the development and auditing of specialized scientific MLLMs, from audio MLLMs to combinations of more than two modalities, such as video, text, audio and brain activity. Using large scale datasets collected by the chair holder and collaborators, GRASP will develop GSP-based methods to ensure that MLLMs are grounded in biological and environmental evidence rather than statistical artifacts. GRASP provides a mathematical bridge between signal processing, deep learning, and neuroscience, offering a GSP-based approach to AI interpretability. The project is supported by the elite computational resources of the GENCI / Jean Zay GPU Cluster and builds upon ongoing PhD tracks, funded by the PRACOM Chair, IMT Atlantique, the Dynabiod PEPR and the EU SEED MSCA project, focused on reasoning faithfulness and MLLM evaluation. By aligning AI interpretability with the communicative principles of the human brain, this project ensures that the next generation of multimodal systems is not only high-performing but fundamentally trustworthy and scientifically grounded.

CALL FOR CHAIRS 2026

AI and strategic reconfigurations for an international-local continuum of governance/ RECONFIAL

Acronym	**RECONFIAL**
Chair Title	**AI and strategic reconfigurations for an international-local continuum of governance**
SequoIA research pilar	☒ Pilar 1: Core AI ☐ Pillar 2: AI, cybersecurity and defense ☐ Pillar 3: AI for environment, ocean
Chair type	☐ Junior ☒ Experienced
Keywords	Interdisciplinarity, public policy, democracy, regulation
Expected beginning / duration	June 2026/4 years
Leading institution	IODE (UMR CNRS 6262)
☒ **Please confirm you have informed your Head of Department of your submission**	

Chair Holder	**Last Name, First Name, Position**	
	TURGIS, Sandrine, Associate Professor of Law (MCF HDR)	
	e-mail	**Phone**
	[redacted]	[redacted]
Co-Chair Holder (if any)	**Last Name, First Name, Position**	
	e-mail	**Phone**

Abstract of the chair project (Non-Confidential – 4000 characters maximum, including spaces). This section will be shared with institutional and industrial partners of the Cluster

In a context marked by two partly correlated phenomena – namely, the acceleration of **AI** capabilities and AI deployment, **on the one hand,** and the proliferation **of hybrid threats and global challenges**, **on the other** – **the significance of the issues of democracy, trust, and security related to AI requires multi-stakeholder and multi-level action**.

Among these, **the international and local levels** are polar opposites, and it might therefore seem counterintuitive to focus on examining their points of convergence regarding AI. Moreover, **the interplay between the international sphere and local authorities** – **particularly with regard to the digital transition and, even more so, the upheavals linked to AI** – has received limited attention in academic literature. **This blind spot in research is compounded by a lack of**

consideration of this dimension on the ground by the actors themselves, with a **few notable exceptions, including the local authorities** that **are institutional partners of the SequoIA Cluster** (Région Bretagne, Rennes metropole, Brest metropole et ville), which position themselves as pioneers. Nevertheless, and quite rightly so, the issue **of the international-local continuum surrounding AI** must find its place within research, legal instruments, and local, national, and international strategies and initiatives in order to holistically address the challenges related to AI and propose effective solutions **for an inclusive, responsible, understandable, and sustainable AI**.

Based on these observations, the Chair "AI and Strategic Reconfigurations for an International-Local Continuum of Governance" (RECONFIAL) will focus on identifying the strategic reconfigurations essential to ensuring an international-local continuum of "AI governance" and of "governing with AI" as a lever for trust and security.

To this end, **the Chair pursues several objectives:**
1. **Map** the territorialization of international issues and the internationalization of local issues driven by AI (= the international-local continuum)
2. **Identify and test** existing approaches to addressing this continuum in "AI governance" and "governance wih AI" (AI territorial charters/strategies, transnational networks, digital diplomacy by local authorities, EU funds for AI transition, etc.)
4. **Propose strategic reconfigurations** of this governance necessitated by the continuum, through combining research, experimentation, and international comparison, in order to inform public policy and help remove structural or cyclical obstacles.

The cross-cutting issues and challenges associated with these strategic reconfigurations extend beyond them, and the chair, **through the synergies** it will foster, **including interdisciplinary ones (social sciences and humanities - law, geopolitics - and digital & computer sciences),** will contribute to **national and international research on the complex challenges posed to society by AI-related transformations**, because IA must be understood not only as a technical object but also as **a major social phenomenon**.

Building on the implementation of a **proven research methodology** that combines systemic studies, cross-cutting approaches, and interdisciplinary dimensions, **the Chair will serve** as **a catalyst** for individual and collective reflection and research, both in France and abroad, with field researches methods involving local actors, international organizations, and state authorities.

Associate professor in international and public law (MCF HDR) at the University of Rennes, member of IODE, **and research associate at CReC**, the Chair holder possesses internationally recognized **legal expertise in artificial intelligence, digital law, and cybersecurity**, **as well as in the internationalization of territories**; strong **integration into** the research **ecosystem**; **and experience** in leading projects, including interdisciplinary ones.

Building on these strengths, this chair will contribute to the visibility of the SequoIA Cluster and its partners.

Cluster SequoIA
Call for chairs
2026

CHAIR PROPOSAL

VISUAL AND INTERACTIVE STORYTELLING FOR AI EXPLAINABILITY/ VISTORIA

Acronym	VISTORIA
Chair Title	Visual and Interactive Storytelling for AI Explainability
SequoIA research pilar	☒ Pilar 1: Core AI ☐ Pillar 2: AI for cybersecurity and defense ☐ Pillar 3: AI for environment and ocean
Chair type	☐ Junior ☒ Experienced
Keywords	Explainable AI, Large Language Models, Storytelling, Human-centered AI
Expected beginning / duration	48 months
Leading institution	Inria
☒ **Please confirm you have informed your Head of Department of your submission**	
Chair Holder	**Last Name, First Name, Position** GALÁRRAGA DEL PRADO Luis, Chargé de Recherche Inria **e-mail** [redacted] **Phone** [redacted]
Co-Chair Holder (if any)	**Last Name, First Name, Position** **e-mail** **Phone**

Summary of the chair project (Non-Confidential – 4000 characters maximum, including spaces)

The ubiquity of black-box AI solutions in many domains of contemporary society calls for the development of transparent and trustworthy models that can provide understandable justifications for their answers to humans. While the explainable AI research community has produced a plethora of solutions to reveal the inner workings of complex AI models, the resulting explanations remain very technical, requiring great effort to be translated into actionable insights understandable to arbitrary audiences. The VISTORIA chair aims to leverage the language modeling and multimodal capabilities of modern large language models (LLMs) to translate existing “technical” AI explanations into understandable and engaging visual interactive narratives. Our approach is human-centered and will first elicit storytelling preferences across different users groups through co-design workshops, thereby informing the narrative generation process to produce user-adapted explanations. These preferences may include rhetorical devices or conceptual metaphors as well as interaction and visual styles. To guarantee faithful non-hallucinatory narratives we will explore different LLM supervision approaches, including prompting, fine-tuning and knowledge fusion techniques grounded on user preferences for storytelling. The resulting techniques and their explanations will then be evaluated through comprehensive user studies that will assess their impact on key cognitive dimensions such as understanding or confidence. Our research aims to contribute to the emergence of a new generation of explainable and trustworthy AI systems that can communicate their inner workings to humans in a pedagogical and natural way through a principled combination of textual narratives and visual assets.

Cluster SequoIA
Call for International Chairs
2026

CHAIR PROPOSAL

SOPHIE PINCHINAT/AI-STAR

Acronym	Sophie Pinchinat/AI-STAR
Chair Title	Artificial Intelligence for Structured Threat Modeling with Attack Trees
SequoIA research pilar	☐ Pilar 1: Core AI ☒ Pillar 2: AI for cybersecurity and defense ☐ Pillar 3: AI for environment and ocean
Chair type	☒ International
Keywords	Risk analysis, Threat modeling, attack generation, attack trees
Expected beginning / duration	September, 2026, 48 months
Leading institution	University of Rennes
☒ **Please confirm you have informed your Head of Department of your submission**	
International Chair Holder	**Lones, Michael, Professor at Heriot-Watt, UK** [redacted]
Local Chair Holder	**Pinchinat, Sophie, Professor at University of Rennes, FR** [redacted]

Cluster SequoIA
Call for International Chairs
2026

CHAIR PROPOSAL

SOPHIE PINCHINAT/AI-STAR

Abstract of the chair project (Non-Confidential – 4000 characters maximum, including spaces)

With advances in computing, our systems have become more ubiquitous. They contribute and impact all aspects of our lives. Recent fast advancing progress in abilities of AI and their use in the development of systems expand the complexity of systems. This leads to a dramatic increase in the security risks associated with these systems. AI although part of problem should also be part of the solution if employed adequately to improve and accelerate security analysis. Risk analysis of multi-domain systems, i.e., made of buildings, secured access, information systems, organizations, etc., together with remediation identification and implementation has become a wide field of research.

The oldest and most widely applied technique in threat modelling are attack trees (ATs) because ATs provide a structured and logical image of a risk and subsequently help organizations better understand the methodologies of their opponents and draws through e.g., quantitative analysis, which prioritized countermeasures to apply. Importantly, ATs can be used effectively in agile environments. Nevertheless, in practice, ATs are build by error-prone manual design or by too informal AI techniques impeding formal reasoning.

The AI-STAR project aims a generic approach for AT-centric threat modelling of multi-domain systems with a strong emphasis on a cross-fertilization between Formal Methods (FM) and AI technologies. It builds on prior collaborations with industry to develop formal tools that highlighted the need for AT-centric formal frameworks aligned with standard threat modeling workflows. However, plain formal tools often have shortcomings: (1) they are inaccessible to practitioners, (2) automated AT generation fails to match expert-designed trees, and (3) AI-based methods produce informal, lacking completeness guarantees.

AI-STAR combines expertise in formal methods (FM) and generative AI (GenIA) to bridge the gap between automated and expert-level assurance artifact generation by using GenAI for creative proposal and FM for verification, ensuring both innovation and correctness.

Sophie Pinchinat's work ranges from Theoretical Computer Science to symbolic AI, including control and supervision. This enables her to bridge these domains effectively, drawing on her deep expertise in formal languages, logic, automata, and game theory, as well as strategic reasoning in multi-agent systems, automated planning, and control problems. Noticeably, her contributions on the formal foundations of ATs have been recognized in leading security conferences.

Michael Lones has an extensive publication record spanning foundational and applied AI that has received international recognition. His work on machine learning pitfalls and safe AI practice is widely cited and has been presented in various academic and practitioner venues. He has significant experience with AI applications to cybersecurity, including malware detection, IoT attack detection, insider-threat identification, and robustness analysis of neural networks.

AI-STAR will create a world-class threat modeling environment with the synergy of scientific expertise and research environments. These complementary domains of the chairs enhance the quality and impact of scientific outcomes that will strengthen the Cluster SequoIA ecosystem's visibility and leadership in large-scale academic and industrial projects.

CALL FOR CHAIRS 2026

CHAIR PROPOSAL
AI-driven Security Against Microarchitectural Vulnerabilities/ AI4MicroSecurity

Acronym	AI4MicroSecurity
Chair Title	AI-driven Security Against Microarchitectural Vulnerabilities
SequoIA research pilar	☒ Pilar 1: Core AI ☒ Pillar 2: AI, cybersecurity and defense ☐ Pillar 3: AI for environment, ocean
Chair type	☒ International
Keywords	Microarchitecture, Hardware Security, Real-time Intrusion Detection, Mitigation, Machine Learning, Secure-by-Design Systems, Trustworthy AI
Expected beginning / duration	September 2026 / 4 years
☒ **Please confirm you have informed your Head of Department of your submission**	
International Chair Holder	**Last Name, First Name, Position** BHATTI, Muhammad Khurram, Associate Professor/Deputy Director Research & Impact (DDoRI) **Institution, Country** University of Exeter, United Kingdom (UK) **e-mail** [redacted]
Local Chair holder	**Last Name, First Name, Position** Gogniat, Guy, Full Professor **Institution** Université Bretagne Sud, Lorient, France **e-mail** [redacted]

Abstract of the chair project (Non-Confidential – 4000 characters maximum, including spaces). This section will be shared with institutional and industrial partners of the Cluster

Chair AI4MicroSecurity will advance AI-driven cybersecurity, focusing on microarchitectural vulnerabilities in modern systems while ensuring AI remains resilient, transparent, and compliant. It aims to reduce exposure to hardware vulnerabilities and strengthen defenses

against Microarchitectural side- and covert-channel attacks (SCAs) that threaten trust, confidentiality, and integrity. The chair supports Cluster SequoIA's priorities, contributing to Pillar 2 (AI, Cybersecurity, and Defense) and Pillar 1 (AI Fundamentals) through research on robust, explainable, and trustworthy AI for security-critical environments.

Recent disclosures of security and privacy vulnerabilities in modern computing systems have revealed deep and systemic weaknesses that span hardware, software, and their interfaces. A particularly severe class of threats is Microarchitectural SCAs, which exploit unintended information leakage arising from the internal organization and optimization mechanisms of processors. Microarchitecture defines how instruction execution, caches, branch predictors, translation lookaside buffers (TLBs), interconnects, and speculative and out-of-order execution units are implemented beneath the Instruction Set Architecture (ISA). While these mechanisms are essential for performance and energy efficiency, they also introduce exploitable leakage channels that are invisible to conventional software-level security models.

The severity of these vulnerabilities is evidenced by a succession of high-profile attacks such as Spectre, Meltdown, Prime+Probe, Flush+Reload, Foreshadow, and Rowhammer. These attacks demonstrate that sensitive data can be inferred via timing, access patterns, and contention in shared components, often bypassing isolation boundaries like processes and virtual machines. Microarchitectural SCAs can be timing-, access-, trace-, or fault-based, and launched locally or remotely. Research shows these threats affect critical national infrastructure (CNI), cloud systems, data centres, edge and AI platforms, and embedded devices.

Chair AI4MicroSecurity will directly address these challenges by strengthening the security foundations of computing systems using advanced AI techniques at the hardware–software interface. The project advances secure-by-design Microarchitectural research to reduce exposure to systemic hardware vulnerabilities and to improve resilience against Microarchitectural SCAs that undermine trust, confidentiality, and system integrity in digital infrastructure. AI-driven systems themselves are susceptible to adversarial manipulation and bias. Therefore, while a major focus of this chair will be to advance AI-based cybersecurity applications, the project will also ensure that AI-driven solutions remain resilient, transparent, and compliant with regulatory frameworks by developing methodologies for detecting and mitigating attacks on AI system itself, such as data poisoning, backdoor exploitation, and misclassification in AI systems.

The chair pursues three integrated aims. First, it will develop AI-enabled, performance-preserving adaptive mitigations through hardware–software co-design, dynamically responding to evolving attack patterns. Second, it will establish a standardized evaluation framework for secure cache architectures, addressing inconsistencies in threat modelling and benchmarking. Third, it will validate these approaches through industrial use cases and offensive AI techniques to support real-world deployment.

Overall, AI4MicroSecurity will enhance resilience to Microarchitectural attacks, strengthen public trust in digital systems, and reduce large-scale cyber risks. It will boost UK–France capabilities in secure processor and firmware design, support skills development, improve supply-chain trust, and enhance competitiveness in semiconductor and cybersecurity sectors, aligning with Cluster SequoIA and national priorities.

CALL FOR CHAIRS 2026

CHAIR PROPOSAL CHAIR NAME / ACRONYM
FEDULCAD

Acronym	FEDULCAD
Chair Title	FEDerated and frUgal Learning for Complex Attack Detection
SequoIA research pilar	☐ Pilar 1: Core AI ☒ Pillar 2: AI, cybersecurity and defense ☐ Pillar 3: AI for environment, ocean
Chair type	☒ International
Keywords	Federated learning, split learning, split federated learning, frugal AI
Expected beginning / duration	10/2026 – 09/2030
☒ **Please confirm you have informed your Head of Department of your submission**	
International Chair Holder	**Last Name, First Name, Position** Omar Abdul Wahab **Institution, Country** **Polytechnique Montréal, Canada** **e-mail** [redacted]
Local Chair holder	**Last Name, First Name, Position** David Espes **Institution** **Université de Bretagne Occidentale, France** **e-mail** [redacted]

Abstract of the chair project (Non-Confidential – 4000 characters maximum, including spaces). This section will be shared with institutional and industrial partners of the Cluster

The FEDULCAD (FEDerated and frUgal Learning for Complex Attack Detection) international research chair aims to develop an innovative Split Federated Learning (SFL) framework for detecting complex cyberattacks in distributed industrial systems (Industrial Control Systems, ICS). As cyberattacks become increasingly sophisticated, coordinated, and difficult to detect, traditional approaches are reaching their limits, particularly due to heterogeneous environments, hardware

constraints, and data confidentiality requirements. SFL emerges as a promising solution by combining the benefits of federated learning and split learning, enabling both knowledge sharing across entities and adaptation to local resource constraints.

However, several major scientific challenges hinder its adoption. First, data heterogeneity (non-IID), which is inherent to industrial systems, introduces complex biases (spatio-temporal and client distribution biases) that degrade the performance of collaborative models. Second, determining the optimal partitioning of the model between client and server is a critical challenge, as it must balance hardware constraints, latency, and detection performance. Third, the security of SFL remains largely underexplored, particularly with respect to new vulnerabilities introduced by dynamic model partitioning.

To address these challenges, the chair is structured around four complementary research axes. The first focuses on developing intelligent model partitioning methods that jointly consider performance and hardware constraints, as well as novel strategies for aggregating heterogeneous models. The second axis aims to optimize collaborative learning mechanisms using reinforcement learning approaches, enabling dynamic adaptation of client participation and improved handling of non-IID data. The third axis is dedicated to analyzing and strengthening SFL security by identifying and mitigating new types of attacks specific to this paradigm. Finally, the fourth axis involves designing a realistic dataset representing industrial systems subject to coordinated attacks, based on an experimental platform that integrates realistic industrial environments.

Led jointly by Polytechnique Montréal and the University of Western Brittany, this chair builds on complementary expertise in federated learning, cybersecurity, and frugal artificial intelligence. It also benefits from strong collaborations with major industrial partners, particularly in the defense and maritime sectors. The project is expected to contribute to the theoretical and methodological foundations of SFL, while enabling tangible advances in cyberattack detection for critical infrastructures.

Beyond its scientific contributions, the chair seeks to structure a Franco-Canadian ecosystem focused on cybersecurity for critical systems, particularly in the maritime domain. It will also contribute to the training of future experts in artificial intelligence and cybersecurity by integrating its outcomes into academic programs at both institutions. Through its interdisciplinary and application-driven approach, FEDULCAD addresses key strategic challenges related to digital sovereignty, infrastructure resilience, and international cooperation.

CALL FOR CHAIRS 2026

CHAIR PROPOSAL CPrivacy-preserving and Trustworthy AI / PYTHAI

Acronym	**PYTHAI (PrivacY-preserving and TrustworthY explainable AI)**
Chair Title	
SequoIA research pilar	☒ Pilar 1: Core AI ☒ Pillar 2: AI, cybersecurity and defense ☐ Pillar 3: AI for environment, ocean
Chair type	☒ International
Keywords	Privacy, Explainability, Trustworthy AI
Expected beginning / duration	September 2026 (4 years)
☒ **Please confirm you have informed your Head of Department of your submission**	
International Chair Holder	**Last Name, First Name, Position** Gambs, Sébastien, Professor **Institution, Country** **Université du Québec à Montréal, Canada** **e-mail** [redacted]
Local Chair holder	**Last Name, First Name, Position** Allard, Tristan, Maître de conférences **Institution** **Université de Rennes** **e-mail** [redacted]

Abstract of the chair project (Non-Confidential – 4000 characters maximum, including spaces). This section will be shared with institutional and industrial partners of the Cluster

PYTHAI (PrivacY-preserving and TrustworthY explainable AI) aims at solving a fundamental tension in modern machine learning: while explanations are increasingly deployed to enhance transparency and trust, they may inadvertently expose sensitive information about the underlying models and their training data. More precisely, the objective of this chair is to systematically study and mitigate the privacy and ethical risks induced by explainability mechanisms, and to design explanation methods that are both informative and robust to adversarial manipulation.

The first research axis focuses on the development of novel privacy attacks that explicitly leverage explanations or model structure. We will investigate how different forms of explanations—such as feature attributions, counterfactual examples, surrogate models or interpretable models—can be

exploited to reconstruct training data or infer sensitive properties (*e.g.*, membership or attributes). This axis aims to characterize which explanation paradigms are inherently more resilient to such attacks and under what conditions.

The second axis, informed by these findings, aims at designing explanation mechanisms that are private by construction. More precisely, we will explore countermeasures that modify how explanations are generated, such as randomized explanation processes to reduce information leakage, or constrained example-based explanations relying on representative subsets ("stereotypes") to limit adversarial gain. Particular attention will be given to complex models such as graph neural networks and to explanation techniques used in graph mining. In addition, we will investigate hybrid approaches combining black-box models with interpretable surrogates trained jointly to balance utility, fidelity and privacy.

The third axis examines how explanations can be manipulated to enable "ethics washing," which can be defined as creating a misleading perception that a model satisfies ethical properties such as fairness or privacy. To achieve this, we will develop methods to quantify the extent to which black-box models can be "fairwashed," for instance by analyzing the variability of fairness metrics within the Rashomon set of high-fidelity surrogate models. This will allow us to derive bounds on achievable fairness and to identify conditions under which such manipulations can be prevented. More broadly, we will investigate other forms of ethics washing, including privacy washing, and propose general mechanisms to detect and mitigate these risks.

Overall, PYTHAI will advance the foundations of safe and trustworthy AI by bridging privacy and explainability. It will lead to principled methods for auditing and designing explanation systems that are resilient to inference attacks and resistant to manipulation. The chair will also foster international collaboration between Brittany and Québec through joint supervision of students and research initiatives that will contribute also directly to the SequoIA ecosystem.