Press Release, 31 March 2026

ALTEN and CentraleSupélec Launch Chair on Trustworthy AI for Critical Industrial Systems

ALTEN, a global leader in engineering and IT services, and CentraleSupélec, France's premier graduate school in science and engineering and a founding member of Université Paris-Saclay, today announced the signing of a chair agreement focused on trustworthy artificial intelligence.

The chair, titled "Trustworthy Multi-Agent Systems for Critical Industries", will advance the scientific and methodological foundations required to deploy explainable, sovereign, and robust AI in high-stakes industrial environments such as energy, transport, manufacturing, and defence. These sectors demand AI systems that are not only high-performing but also transparent, auditable, and resilient in the face of operational and regulatory challenges.

Professor Wassila Ouerdane, a computer science expert at CentraleSupélec and a member of the MICS Laboratory (Mathematics and Computer Science for Complex Systems), will lead the chair. She will oversee public-private research collaborations, with a scientific launch conference planned for autumn 2026.

Funded initially for three years and supported by the Fondation CentraleSupélec, the chair builds on a long-standing partnership between ALTEN and CentraleSupélec. Over the past five years, the two institutions have collaborated on engineering education and joint PhD research, creating a seamless continuum of training, research, and innovation. This initiative will further intensify research in strategic areas, including digital sovereignty and the robustness of vital French industrial systems.

The chair will directly inform ALTEN's research on critical multi-agent systems, with scientific results widely published to benefit the broader ecosystem including industry professionals, researchers, and students involved in developing these systems.

Additionally, the chair aims to progressively engage industrial partners from manufacturing, energy, transport, and defence who seek to accelerate the adoption of multi-agent systems in critical contexts, all while reinforcing trust through a shared scientific approach.

Yann Rougemont, Innovation Director at ALTEN, stated:
"With this chair, we are taking a significant step forward in achieving our goals: providing industrial clients with genuinely reliable AI, embedded at the core of their most critical operations. This translates into multi-agent systems capable of explaining their decisions, detecting errors, and collaborating securely in complex environments whether in factories, energy networks, or fleets of autonomous vehicles.

By leveraging CentraleSupélec's cutting-edge research, ALTEN will transform scientific breakthroughs into operational solutions: real-time optimisation, reduced unplanned downtime, enhanced resilience, and native compliance with the AI Act. This is trustworthy AI transparent, robust, and sovereign. "

Paul-Henry Cournède, Research Director at CentraleSupélec, added:
" The chair established with ALTEN is the 8th AI-focused chair at our institution, reinforcing our position as France's leading school for AI education and research within the unique ecosystem of Université Paris-Saclay. It addresses a highly strategic topic trustworthy AI and highlights the excellence of our research teams in supporting French industries as they tackle challenges related to the reliability and robustness of deployed infrastructures and systems.

Sensitivity: C1-Internal

Autonomous multi-agent systems raise novel scientific questions that traditional explainability and verification methods cannot resolve. This chair provides the resources to address these fundamental issues using real industrial data and constraints, which is precisely what the scientific community needs to generate impactful results."

The Rise of Multi-Agent Systems in Modern Industry: A Scientific Challenge

Modern industry faces growing operational complexity, real-time adaptability, and resource optimisation challenges. In this context, multi-agent systems (MAS) are emerging as an essential scientific and technological solution, transforming traditional industrial processes into intelligent, decentralised, and resilient ecosystems.

However, deploying these technologies in critical environments raises fundamental questions that academia and industry cannot address in isolation:

How can AI-driven agent systems ensure traceable, auditable decisions?

How can autonomous agents collaborate safely without undesirable emergent behaviours?

How can regulatory requirements (e.g., AI Act, GDPR) be integrated from the design phase rather than as an afterthought?

The ALTEN-CentraleSupélec Chair will systematically explore these challenges, aiming to develop trustworthy AI solutions for the most demanding industrial sectors.

"Trustworthy AI" is built on three complementary pillars: lawfulness (compliance with current regulations), robustness (the ability to handle errors, approximations, and uncertainties), and ethics (adherence to human values and principles).

In the context of distributed industrial systems, these requirements translate into concrete scientific challenges explaining collective decision-making processes, formalising and integrating industry-specific constraints, and controlling error propagation across Cloud, Edge, and Mist computing architectures.

Scientific Objectives and Ambition

The chair has a core objective: to establish the scientific, technological, and methodological foundations for developing distributed, robust, and explainable systems based on multi-agent architectures. These systems must be capable of operating autonomously and in a coordinated manner within critical industrial environments.

Here, "criticality" is understood in its broadest sense, encompassing the safety of people and infrastructure, technological and decision-making sovereignty, data quality and governance, computational efficiency, and human integration in decision-making processes. These technical, strategic, and societal dimensions form the framework of requirements that the research must address.

The research conducted under this chair will focus on the following overarching question: How can we design trustworthy distributed intelligent agent systems that operate in critical environments while remaining understandable, verifiable, and compliant with user and regulatory requirements?

Three Complementary Research Axes

To achieve its ambitious goals, the chair's work is structured around three closely linked research themes, each addressing a fundamental dimension of trustworthy AI in distributed systems.

Sensitivity: C1-Internal

Theme 1: Explainability and Trust in Multi-Agent Systems

This theme tackles the core challenge of decision intelligibility in distributed architectures. While explaining a single AI model is already complex, the task becomes exponentially more difficult when global behavior emerges from interactions among multiple autonomous agents. Local explanations from each agent must be composed into a coherent, auditable justification at the system level.

The research focuses on four complementary axes:

1. Multi-level explainability models to link local agent decisions with collective system dynamics.
2. Causal explanation and systemic traceability, using causal models adapted to distributed architectures.
3. Human-centered interactive explainability mechanisms, accounting for diverse user profiles and cognitive loads.
4. Metrics and evaluation protocols to assess the quality and fidelity of generated explanations.

Theme 2: Requirement-Driven Hybrid AI for Complex Systems

This theme forms the methodological backbone of the chair. It starts from the observation that in critical systems, predictive performance alone is insufficient. Systems must also meet explicit requirements for safety, robustness, frugality, and sovereignty requirements that must shape the entire design cycle, from internal representations to decision-making mechanisms.

The approach integrates three types of constraints into hybrid architectures combining statistical learning and symbolic knowledge:

- Structural constraints: Embedding business rules, operational invariants, and safety levels directly into representation and learning mechanisms.
- Dynamic constraints: Using Physics-Informed AI to ensure decisions align with physical laws and real system dynamics.
- Global behavioral constraints: Translating critical requirements into verifiable formal properties for individual and collective agent behaviors

Theme 3: Sovereignty, Robustness, Reliability, and Security in Distributed Multi-Agent Systems

This theme addresses specific risks that arise when AI systems are deployed in distributed Cloud, Edge, and Mist computing architectures. It focuses on three categories of operational reliability threats:

1. Hallucinations in generative models embedded in agents.
2. Error propagation between agents.
3. Threats to factuality and data integrity in exchanged information.

The research explores :

- Real-time hallucination detection mechanisms, combining symbolic and machine learning approaches, and dynamic correction strategies to request human or external validation when confidence levels are low.
- Error propagation modeling in distributed architectures and containment protocols to prevent cascading effects.

Sensitivity: C1-Internal

- Factchecking and data integrity verification mechanisms, using traceability and security techniques tailored to Edge and Mist computing constraints.

High-Impact Applications in Critical Industries

The chair's research is designed to deliver direct, practical applications in four high-stakes industrial sectors, where trust in multi-agent systems is essential:

1. Industrial Sector: Trust = Performance + Traceability

In smart factories, distributed multi-agent systems manage automated production, predictive maintenance, and real-time logistics optimization, ensuring efficiency in normal operations and resilience during crises.

ALTEN Use Case:
A connected factory uses Edge AI agents to monitor and optimize automated production in real time, while coordinating with Cloud AI agents that manage demand and supply chain logistics.

2. Energy Sector: Trust = Resilience + Infrastructure Transparency

In smart grids and decentralized energy networks, distributed agents handle energy production, storage, and distribution.

ALTEN Use Case:

A smart grid combines strict symbolic rules for managing peak demand with data-driven optimization algorithms to adjust production in real time.

3. Transport Sector: Trust = User Safety + Accountability

Distributed multi-agent systems are transforming mobility, from autonomous vehicles to smart urban traffic management and connected logistics.

ALTEN Use Case:

In an urban traffic management system, explainable AI clarifies why an agent prioritized one traffic flow over another, preventing congestion or accidents caused by unclear decisions. Similarly, a Cloud/Edge multi-agent system enables autonomous vehicle fleets to share real-time data (e.g., road conditions, obstacles) to optimize last-mile delivery.

4. Defence Sector: Trust = Operational Safety + Responsibility

Distributed multi-agent systems (e.g., drones, robotic swarms, electronic warfare sensor networks) are critical to modern military operations, where autonomy and coordination are vital.

ALTEN Use Case:

A drone swarm integrates:

- Explainability to justify tactical decisions,
- Hybrid AI to comply with engagement rules while adapting to terrain,
- Sovereignty and robustness to operate securely in hostile environments.

Sensitivity: C1-Internal

About ALTEN – www.alten.com

Founded in 1988, ALTEN is a global leader in engineering and IT services, supporting clients in innovation, R&D, and technological IT systems. The company works with key players in aerospace, defense, security, naval, automotive, rail, mobility, energy, environment, life sciences, industrial equipment, electronics, telecoms, banking, finance, insurance, retail, media, and public services.

ALTEN employs over 57,000 people across more than 30 countries, with 88% of its workforce being engineers. In 2025, the group reported €4.1 billion in revenue.

Press Contact: ALTEN Press Office – alten@hopscotch.fr

About CentraleSupélec – www.centralesupelec.fr

Born in 2015 from the merger of École Centrale Paris (1829) and Supélec (1894), CentraleSupélec is one of France's leading graduate schools in engineering and systems sciences. As a founding member of Université Paris-Saclay, it trains over 5,400 students annually across four campuses in Paris-Saclay, Metz, Rennes, and Reims, and boasts a global alumni network.

With 19 research laboratories, CentraleSupélec thrives in a unique innovation ecosystem, partnering with 3,300+ companies and 1,270+ startups, including 10 unicorns with a combined valuation of €82 billion. The school welcomes 25% international students and collaborates with 170+ prestigious universities worldwide. As president of the Groupe des Écoles Centrale, it also oversees campuses in Beijing, Hyderabad, and Casablanca.

Driven by an ambitious strategic plan, CentraleSupélec aims to double its number of graduates between 2022 and 2032 and strengthen its impact on major societal challenges, including ecological and energy transitions, national and European sovereignty, health, and quality of life.

About Fondation CentraleSupélec – www.fondation-centralesupelec.fr

The Fondation CentraleSupélec has a dual mission: to enhance the international reputation of CentraleSupélec's higher education and scientific research, and to support students, faculty, and researchers in funding and developing their projects.

As one of France's largest higher education foundations, it brings together 7,000 individual donors and over 100 corporate sponsors, raising an average of €9 million annually.

Press Contacts:

- Claire Flin – clairefline@gmail.com | +33 6 95 41 95 90
- Marion Molina – marionmolinapro@gmail.com | +33 6 29 11 52 08

Sensitivity: C1-Internal