Spring 2026 Newsletter

faith-ec-project.eu

Spring 2026

Follow Us

Follow us on social media for the latest updates, exclusive content, and more!

In this issue

In this edition, we explore key topics related to the **FAITH** project, including its latest developments, insights, and innovations. Below is the full index of this issue:

Funded by the European Union under GA No 101135932. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union. Neither the European Union nor the granting authority can be held responsible for them.

Tachos Nikos

Senior researcher

FIEEE, FEAMBES, FIAMBE

Department of Materials Science and Engineering of University of Ioannina and in the Foundation for Research and Technology –Hellas / Biomedical Research Institute (FORTH/BRI)

[ntachos@bri.forth.gr]

FAITH Enters Operational Phase: Trustworthy AI Validated Across Seven Large-Scale Pilots

During the second phase of the **FAITH project**, substantial progress has been achieved as the consortium transitions from framework preparation into **full-scale operational deployment across the seven Large-Scale Pilots.** Building on strong foundations established during the first phase, FAITH has entered a highly productive stage focused on **real-world validation, cross-domain experimentation**, and **refinement of its human-centric trustworthiness ecosystem**. The project remains fully aligned with its Description of Action and amended implementation schedule, with milestones and activities progressing according to plan.

A central milestone of this phase is the operational launch of the first integrated FAITH software suite supporting trustworthiness assessment in practical environments. This includes the deployment of **TrustGuard**, the core orchestration platform implementing the FAITH AI Trustworthiness Assessment Framework (AI_TAF), together with **TrustSense,** the maturity and readiness assessment tool for AI teams and stakeholdersas well as an **AI Model Hub**, which incorporates the AI Model Passport and Data Passport functionalities.

These interoperable tools now provide the consortium and pilot owners with a unified environment for analysing risks, documenting AI assets, improving transparency, and strengthening accountability throughout the AI lifecycle. An important achievement of this period is that the seven Large-Scale Pilots have now started actively **using the FAITH AI_TAF supporting tools**, as part of their operational workflows and **trustworthiness assessment** activities. This marks a decisive step from methodology design to practical implementation, allowing pilot partners to evaluate their AI systems under real conditions while systematically documenting assets, identifying vulnerabilities, assessing human readiness, and applying mitigation measures. The practical use of these tools across heterogeneous sectors demonstrates the adaptability and scalability of the FAITH ecosystem.

The seven Large-Scale Pilots are now entering or expanding their execution phases across highly relevant domains: media, transportation, education, underwater robotics, industrial wastewater management, healthcare, and active ageing. In each

trustSense
Human evaluation of trustworthiness

Trust Guard

case, domain-specific AI solutions are being tested under **real operational conditions** while simultaneously applying the FAITH trustworthiness methodology. This combination of technical deployment and trust evaluation is one of the project's most distinctive strengths. Pilot partners are generating valuable **evidence on fairness, usability, privacy, explainability, robustness, and human oversight,** while also identifying sector-specific requirements that will further enhance the FAITH framework.

Strong progress has also been achieved in cross-pilot collaboration and knowledge transfer. Through dedicated **workshops, dry-run exercises, bilateral technical meetings,** and **plenary events**, the consortium has established a dynamic mechanism for sharing lessons learned among domains. This ensures that innovations developed in one pilot can inform others accelerating maturity and maximizing impact.

The project's multidisciplinary nature, combining technical, legal, ethical, societal, and human factors expertise, continues to be a major asset in delivering a practical and scalable trustworthy AI solution. Dissemination, communication, and stakeholder engagement activities have continued to expand significantly. FAITH has strengthened its visibility through scientific publications, conference presentations, workshops, webinars, and a continuously growing online presence. Engagement with policymakers, industry, researchers, and end-users is helping position FAITH as an important European reference initiative in trustworthy and human-centric AI. At the same time, exploitation planning and market monitoring activities are advancing to ensure long-term sustainability of the project's key exploitable results.

Overall, FAITH is progressing nicely and is now demonstrating the concrete value of its methodology in r**eal-world critical domains**. The project has moved decisively from conceptual design into operational validation, demonstrating excellent potential to deliver **impactful scientific, technological, societal,** and **policy outcomes** during the next stages. By combining innovation with responsibility, FAITH continues to pave the way for trustworthy AI systems that genuinely serve people, organisations, and society.

LSP3 – Education: Trustworthy AI for Deeper Learning in STEM Education

Prof. Gregoris Mentzas
Founder and Director of the Information Management Unit (IMU) at ICCS, [gmentzas@mail.ntua.gr]
Dr. Dimitris Koulentianos
Physicist, Research & Development Department, Ellinogermaniki Agogi, [dkoulentianos@ea.gr]

The Education LSP, one of the seven Large Scale Pilots (LSPs) of the **FAITH project**, focuses on assessing the **trustworthiness** of novel, AI-based assessment methods of students' complex skills as problem-solving. This is realized through inquiry-based activities for STEM subjects, ensuring that the **AI tools shaping the next generation of learners** are not only effective, but genuinely trustworthy.

The Educational LSP AI Solution, available through **trust-ai-lab.eu**, offers AI-enhanced **personalized learning** and **formative student assessment**. It also allows students to pursue goals that require extended engagement or persistence across multiple contexts and learning opportunities and also enabling teachers to track student progress.

Since the launch of the 1st Pilot on May 2024, the solution developed by **IMU/ICCS** has been deployed extensively. Inquiry-based scenarios have been implemented repeatedly at **Ellinogermaniki Agogi** and in other schools across Greece and other European countries, serving both to **validate** the pedagogical approach and to **train the AI models** of the platform.

Alongside student implementations, a series of workshops and teacher training activities have been delivered. During these sessions, teachers explored the platform, engaged with the technical team on the AI methodology behind the solution, and took part in discussions on **trustworthy AI in education**. Participants then completed questionnaires evaluating the trustworthiness of the AI tools in the context of inquiry-based learning, directly feeding into the project's impact assessment.

From the outset, the solution has been built around a clear vision: that trustworthy AI and AI should be synonymous. This means that every design and implementation decision is guided by the core trustworthiness dimensions most relevant to educational AI:

- **Validity**, thus ensuring AI-generated content is scientifically accurate and pedagogically sound, directly impacting student engagement and learning outcomes.
- **Explainability**, thus making the AI's reasoning transparent and understandable to both teachers and students, fostering trust and informed use of the system.
- **Fairness & non-discrimination**, which is designing assessments and recommendations that are unbiased and equitable across all learners.
- **Security & Data Sovereignty**, which is protecting student data and ensuring compliance with privacy and security standards throughout the platform.

A rigorous in-isolation trustworthiness assessment of the Educational LSP AI Solution has been completed across all AI lifecycle phases, using the **FAITH-AI - TAF tools**. The impact assessment identified validity and explainability as the most critical dimensions of trustworthiness for the education LSP.

To conduct the trustworthiness assessment, ICSS and EA participated in dedicated training workshops on the use of the FAITH-AI-TAF tools. The Replication Pilot is now underway, with plans to conduct an in-context trustworthiness assessment of the solution in **real classroom environments across Europe.**

The Educational LSP AI Solution represents a new standard for responsible AI in education, one that places student learning, teacher empowerment, and trustworthiness at its heart.

TrustGuard: Evaluating Trustworthiness in AI Systems

Dr. Eleni Tsalapati
Research Project Manager, ATC S.A. [etsalapati@gmail.com]

TrustGuard is an in-isolation AI risk assessment system designed to evaluate risks arising from the compromise of key trustworthiness characteristics—such as **safety, security** and **resilience**—across both model and data assets throughout the entire AI lifecycle, from design and development to deployment. In addition to identifying risks, the tool provides asset-level controls to support targeted mitigation.

Built as a checklist-based solution, TrustGuard operationalises **a six-phase, risk-based methodology** defined within the FAITH trustworthiness assessment framework and is aligned with ISO 27005 and ISO 42001 standards.

The tool assesses trustworthiness dimensions derived from leading European and international best practices and standards, including guidance from the EU High-Level Expert Group on AI (HLEG), ENISA, NIST's AI Risk Management Framework, and relevant ISO/IEC initiatives. Together, these references form a comprehensive framework for evaluating and promoting trustworthy AI systems.

TrustGuard begins with the cartography phase, which defines the assessment boundaries—such as the AI system assets, the lifecycle stage under evaluation, the relevant stakeholders, and the system's sector-specific criticality, in line with the **AI Act**.

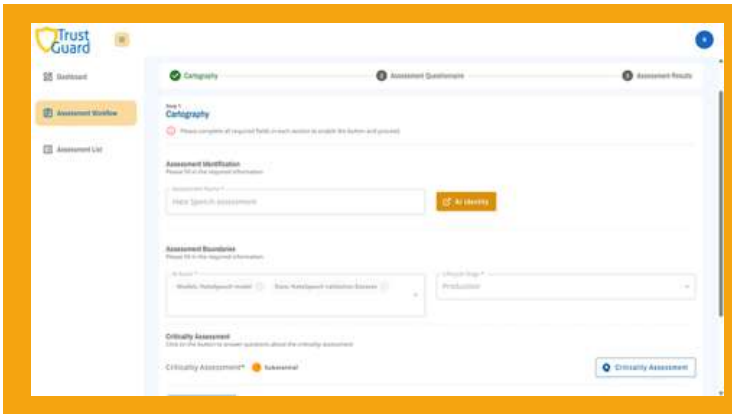

It then moves to the AI trustworthiness **threat assessment phase,** where the user evaluates potential technical threats to the AI system, (e.g., data poisoning, evasion) based on their sector-specific likelihood.

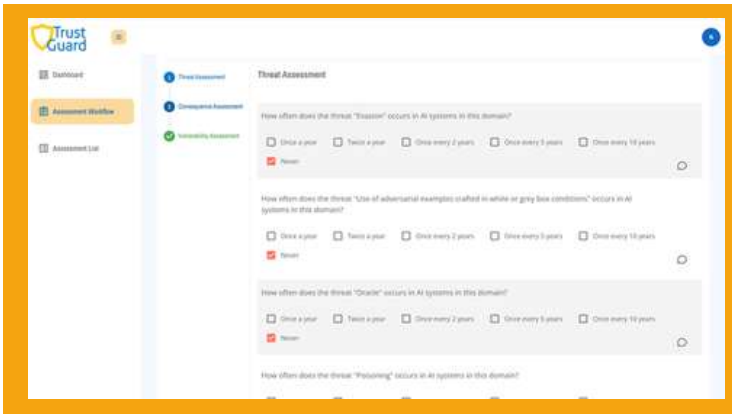

The third phase focuses on assessing the sector-specific severity of potential **business, ethical, or legal consequences** that may arise from a lack of certain trustworthiness properties - consequences that could result if specific technical threats materialize.

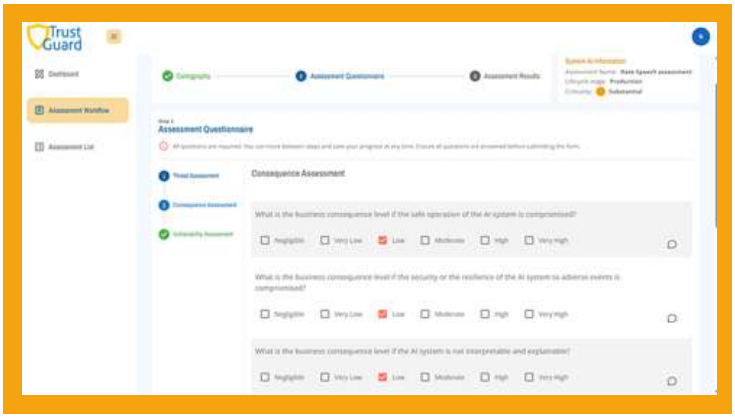

In phase four, the analysis shifts to system vulnerabilities. Here, the user answers a set of technical control questions (e.g., “Is the training dataset enriched with adversarial examples?”) to determine AI system vulnerabilities to AI threats.

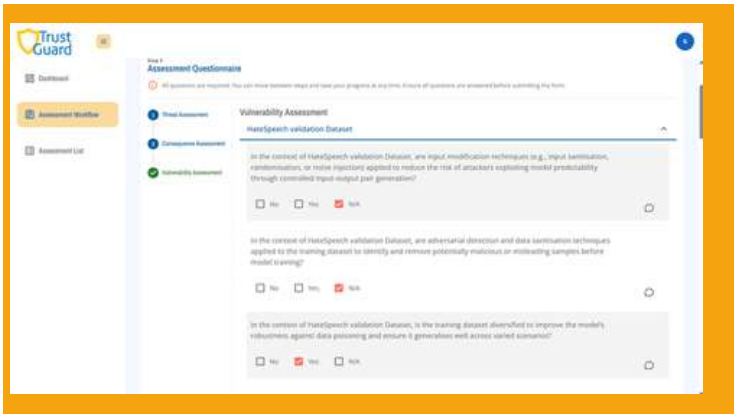

Phase five involves calculating risk levels by combining the likelihood and impact of identified threats, based on the context-specific user’s inputs. Finally, phase six presents the threatened trustworthiness properties, the associated technical threats and risks, and a corresponding set of mitigation actions. These include technical safeguards, governance mechanisms, and behavioural or social interventions.

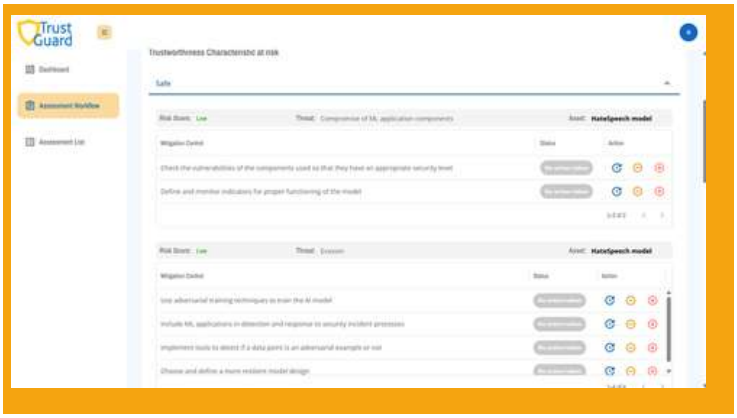

This end-to-end approach captures the complexity and interdependence of trust-related risks, ensuring that each assessment is context-specific and aligned with the unique characteristics of the AI system under test.

The current prototype version of TrustGuard is now being tested by the project’s Large-Scale Pilots (LSPs). The feedback collected during this phase will directly inform the development of a second, more mature version of the system.

FAITH in Practice: AI, Robotics and Port Innovation

An interview with Andrea Minardi, Smart Port and Security Manager at Port of Ravenna Authority *[andrea.minardi@port.ravenna.it]*

The Port of Ravenna has been actively involved in the FAITH project. Can you briefly explain what has been done so far?

Within the FAITH project, the Port of Ravenna has hosted one of the large-scale pilots aimed at testing the trustworthiness of artificial intelligence systems applied to robotics, and specifically the **underwater robotics**. The activities focused on the use of autonomous underwater vehicles equipped with sensors and AI algorithms for two main use cases: automatic recognition of submerged objects and autonomous coverage for bathymetric surveys. In addition, we organised two stakeholder workshop in Ravenna, held in hybrid format, which brought together around 40 participants including port operators, technical experts and institutional stakeholders . These were key moments to present the pilot activities, discuss risks related to AI adoption and collect direct feedback from end users.

Ravenna also hosted a demonstration in April 2026. What was the purpose of this demo?

Yes, in April 2026 we hosted the live demonstration of the FAITH pilot in the Port of Ravenna. The demo represented a crucial step in validating the technologies in a real operational environment. It allowed us to test the performance of AI-based systems integrated with **underwater drones in real port conditions**, particularly in complex and potentially hazardous scenarios. The demonstration showed how these systems can autonomously collect data, detect anomalies and support decision-making processes related to infrastructure monitoring.

What were the main outcomes emerging from the stakeholder engagement activities?

The stakeholder workshop played a fundamental role. It was not just a dissemination event, but a true co-creation exercise. Participants actively contributed by identifying gaps in the proposed systems and by defining priorities for future improvements . One of the key outcomes was the confirmation that trust, transparency and risk mitigation are essential conditions for the adoption of AI in port environments. The discussion also highlighted the importance of aligning technological development with operational needs and regulatory requirements.

What were the results so far from the FAITH project implementation?

The pilot has already delivered several important results. First, it enabled the collection of real-world underwater data, which has been used to train and improve AI algorithms for object recognition and navigation .

Moreover, the tools have been adapted following the suggestions coming from the stakeholders in terms of usability, accountability, security, and first of all, the possibility to take control of the vehicle at any moment and modify the mission.

Andrea Minardi

Andrea Minardi is the Smart Port and Security Manager at Port of Ravenna Authority

Partner Presentation

Greece
iccs.gr

ICCS:
Bringing safe, effective, and human-centred AI into education

The **Institute of Communication and Computer Systems** is one of Greece's premier research organizations, situated within the National Technical University of Athens (NTUA). Structured around a constellation of interdisciplinary research groups and laboratories, ICCS brings together academic staff, researchers, PhD students, and faculty members from NTUA to address pressing scientific, engineering, and societal challenges.

Among ICCS's research units, the Information Management Unit (IMU, imu.ntua.gr) focuses on advanced information management, artificial intelligence (AI), and trustworthy digital innovation. Since its founding in 1997, IMU has evolved into a multidisciplinary research unit focused on designing intelligent, reliable, and human-centered AI services and advanced technologies for digital transformation. The unit's research portfolio encompasses the entire data and machine intelligence lifecycle, from acquisition and analysis, through secure computation, to decision support. A central feature of IMU's identity and vision is a strong commitment to trustworthy AI. The IMU's Trustworthy AI Lab (http://imu.ntua.gr/trail/) was established specifically to empower the emergence of human-centered, reliable, and responsible AI.

The lab's explicit mission is to make "trustworthy AI" synonymous with "AI", ensuring that all systems are designed to mitigate bias, guarantee transparency and security, uphold data sovereignty, and protect against discrimination.

Within the FAITH project, ICCS/IMU has led two critical technical workstreams:

- Development of the Educational LSP AI Solution: ICCS designed and developed the Educational LSP AI solution at the heart of the FAITH project, enabling innovative AI-driven learning experiences for students.
- Trustworthiness Assessment: ICCS conducted trustworthiness assessment of the Educational LSP solution using the FAITH-AI TAF tools, evaluating the system against criteria of fairness, transparency, security, and reliability.

Through its foundational technical contributions and its leadership in trustworthy AI, ICCS/IMU stands as a key pillar of the FAITH project's mission to bring safe, effective, and human-centred AI into education.

ELLINOGERMANIKI AGOGI

Greece
ea.gr

Ellinogermaniki Agogi (EA): Leading European partner in technology-enhanced education

Ellinogermaniki Agogi (EA) is one of Europe's most innovative schools, with over 60 years of operation. As institutional member of MINT EC and UNESCO, EA combines educational excellence with a long-standing commitment to research, technological innovation, and international collaboration.

Since 2002, EA has hosted a certified Centre for Teachers' Continuing Professional Development, training over 800 teachers annually. The school's Research and Development Department bridges pedagogical research and technological innovation, offering annual Technology-Enhanced Education courses aligned with Ministry of Education guidelines. To date, EA has participated in more than 300 EU-funded projects as partner or coordinator.

Within the FAITH project, EA has taken on a central and multifaceted role across three key areas:

- Teacher Training: EA organised workshops, summer schools, and professional development activities for teachers, building educator capacity across the project partnership.
- Student Implementations: EA conducted over 1,200 in-class implementations with students using the educational LSP AI solution, representing one of the largest deployment records within the project.
- Impact Assessment: EA led the impact assessment of the educational LSP AI solution, generating robust evidence on its effectiveness and informing its future development and scalability.

Through this work, EA continues to demonstrate its role as a leading European partner in technology-enhanced education and evidence-based innovation.

Stakeholders event

Workshop in Ravenna Showcases Advances in Trustworthy AI-Driven Underwater Robotics

The FAITH Project recently held a successful stakeholder workshop at the port of Ravenna, bringing together project partners, researchers, technology experts, and end users to discuss the latest progress in trustworthy AI-driven underwater robotics.

The workshop focused on two core technological pillars supporting the project's robotics pilot activities: Automatic Target Recognition (ATR) and Autonomous Coverage (AC). These advanced capabilities are being developed to enhance the performance of autonomous underwater vehicles in detecting and geolocating submerged objects and anomalies, while also enabling adaptive bathymetric surveying for port infrastructure inspection and monitoring.

Participants engaged in productive discussions on the operational potential of these technologies and their role in improving efficiency, reliability, and safety in underwater inspection missions.

A major highlight of the event was the live demonstration of the "Hyppo" prototype, developed for underwater inspection and data acquisition activities. The demonstration provided attendees with a firsthand look at the system's capabilities in real operational conditions and showcased the project's ongoing efforts to deliver innovative and trustworthy AI-enabled robotic solutions for maritime environments.

The workshop further strengthened collaboration between research and industry stakeholders, reinforcing the project's commitment to ensuring that technological innovation remains closely aligned with real-world operational needs.

Through events such as this, the FAITH Project continues to advance the development of reliable AI technologies for safer, smarter, and more efficient underwater operations.

Publications

- Maglaras, L., Kioskli, K. **End-to-End Encryption: Technological and Human Factor Perspectives** In: Schmorrow, D.D., Fidopiastis, C.M. (eds) Augmented Cognition. HCII 2025. Lecture Notes in Computer Science(), vol 15778. Springer, Cham. https://doi.org/10.1007/978-3-031-93724-8_10
- Topini, Alberto & Cecchi, Lorenzo & Fedi, Fausto & Minarelli, Marco & Bucci, Alessandro & Ridolfi, Alessandro **Trustworthy AI-Driven Autonomous Underwater Vehicles for Port Infrastructure Inspection: Paradigm Conceptualization** 2025 IEEE International Workshop on Metrology for the Sea; Learning to Measure Sea Health Parameters (MetroSea) https://10.1109/MetroSea66681.2025.11245714
- Dr Theofanis Fotis , Kitty Kioskli, Eleni Seralidou **Charting Trustworthiness: A Socio-Technical Perspective on AI and Human Factors** Human Factors in Cybersecurity. AHFE (2025) International Conference. AHFE Open Access, vol 168. AHFE International, USA http://doi.org/10.54941/ahfe1006137

More publications

@faith_ec_project