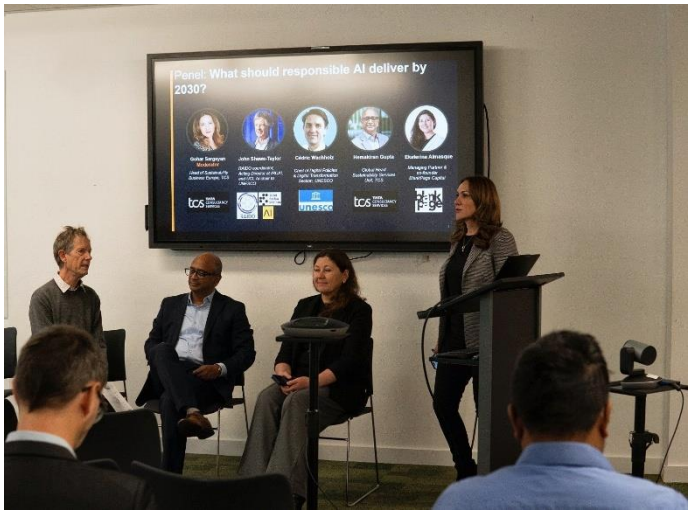
RAIDO

PRESS RELEASE
November 2026

RAIDO Presents Its Final Results for Trustworthy and Green AI in London



The RAIDO project presented the results of three years of research and innovation at its Final Industry Showcase, held on 16 November 2026 at the UCL Centre for Artificial Intelligence in London, United Kingdom. Under the title *"AI That Works for People and Planet"*, the event combined the RAIDO Final Industry Showcase with the UCL UNESCO Chair on AI Forum, highlighting how trustworthy and sustainable Artificial Intelligence can address real-world challenges and support the needs of industry and society. The event was convened by Professor John Shawe-Taylor, RAIDO's Principal Investigator at IRCAI (the project's coordinating organisation), together with RAIDO partners and the UCL AI community.

The morning session focused on the transition from European research to practical AI solutions, providing an opportunity for industry representatives, researchers and potential users to explore the technologies developed within RAIDO. A central objective was to demonstrate how organisations can benefit from more resource-efficient, transparent and reliable AI systems, while exploring opportunities for the wider adoption and future use of the project's results.



Through presentations, demonstrations and a dedicated poster session, the RAIDO consortium presented its integrated AI platform, the outcomes of its four real-world pilots in energy, agriculture, healthcare and robotics, and the project's Key Exploitable Results (KERs). Particular emphasis was placed on measurable improvements in data efficiency, energy consumption, model performance and AI trustworthiness, illustrating the practical value of the developed technologies for end-users across different application domains.



From Research to Practical AI Solutions

During the morning session, MetaMind Innovations (MINDS) introduced the project's technological approach through the presentation *"RAIDO in plain English: what we built and why it matters"*. The presentation addressed three major challenges affecting the adoption of AI in real-world environments: the limited availability of high-quality training data, the energy and computational demands of AI models, and the need for greater transparency and trust in AI-driven decisions. To address these challenges, RAIDO has brought together a range of technologies within a single integrated platform, designed to support the development and deployment of reliable, energy-efficient and trustworthy AI solutions.

The RAIDO platform brings together 13 integrated components through a single web interface, supporting the entire AI lifecycle, from data preparation and model training to deployment and performance monitoring. Its capabilities include automated data annotation, synthetic data generation, energy-efficient model training, explainable AI and blockchain-based traceability. By combining these technologies, the platform enables developers and organisations to make better use of available data, reduce computational requirements, monitor energy consumption and gain a clearer understanding of AI-generated results.

As part of the session *"RAIDO Journey: Real-World Stories"*, Christos Kotronis and Giorgos Theodorou from UBITECH introduced the platform's core functionalities through a video demonstration aimed at a business and industry audience. The session continued with a presentation by Dr Olga Segou from Netcompany, focusing on how the RAIDO technologies have been applied, tested and validated across different real-world scenarios, demonstrating their potential to address specific operational needs and deliver measurable benefits to end-users.



Real-World Applications and Pilot Results

A central part of the Final Industry Showcase was the presentation of RAIDO's pilot activities, which demonstrated how a shared AI architecture can address different challenges across four application domains: energy, precision agriculture, healthcare and robotics. By testing and validating its technologies in different operational contexts, RAIDO explored how organisations and end-users can benefit from more efficient use of data and computing resources, improved AI model performance, greater transparency and reduced environmental impact.

Pilot 1 – Energy Domain

In the energy domain, RAIDO addressed challenges related to photovoltaic generation forecasting and energy management across distributed smart home environments. Using real photovoltaic data from the AYESA headquarters in Seville, the project achieved a 66% reduction in AI model size while maintaining 97% fidelity to the original teacher model. In a second use case, using CERTH's Smart Home data to simulate a distributed network of households, RAIDO achieved a 43% reduction in CO_2 emissions, a 72.8% reduction in model size and a 21.2% improvement in model fairness. These results demonstrate the potential of RAIDO's technologies to help energy providers and smart energy system operators develop more resource-efficient forecasting models while preserving the privacy of household data.

Pilot 2 – Precision Agriculture Domain

In the precision agriculture domain, RAIDO worked with Biolynx to improve the monitoring of fungal fermentation processes for sustainable food production. For early contamination detection, the project achieved a 32% reduction in model size, using 50% less data and consuming 22.3% less energy, while improving accuracy by 3%. In fungal growth prediction, the use of up to 99.6% synthetic data resulted in a 3% increase in accuracy. A third use case, focusing on fungal morphology classification, required 30% fewer labelled images and achieved 23.9% shorter training time and 29.9% lower emissions. These results highlight the potential of RAIDO's technologies to reduce the need for physical sampling and manual annotation, minimise contamination risks and support more efficient and sustainable food production processes.

Pilot 3 – Healthcare Domain

In the healthcare domain, RAIDO focused on supporting clinical decision-making while addressing challenges related to data privacy, fairness and trust in AI. In collaboration with Jessa Hospital, the project used synthetic patient case data to improve an AI-powered pharmacogenomics (PGx) support system, achieving a 62% increase in trust in AI-generated PGx findings and an 80% acceptance rate. In a second use case, focusing on cancer relapse prediction, federated learning enabled collaborative model training without transferring sensitive patient data between clinics, achieving 99.19% accuracy and only 0.5% performance disparity across participating clients. These results highlight the potential of RAIDO's technologies to support healthcare professionals in accessing more reliable AI-driven insights while respecting patient privacy and data governance requirements.

Pilot 4 – Robotics Domain

In the robotics domain, RAIDO explored how AI can accelerate the mechanical characterisation of plant fibres for the development of sustainable bio-based composite materials. By combining synthetic data generation, model training and explainable AI, the project developed an alternative to computationally expensive simulations, achieving 300 times faster predictions with over 95% accuracy. In a second use case involving microrobotic testing, RAIDO achieved 480 times faster mechanical property predictions with 97.7% accuracy. These results demonstrate the potential of RAIDO's technologies to help researchers, engineers and materials developers reduce computational costs, accelerate materials testing and gain a better understanding of the factors influencing material properties.

Together, the four pilots demonstrated the versatility of the RAIDO platform and its ability to address common AI challenges across different sectors. By reducing dependence on large volumes of real-world training data, improving computational and energy efficiency, and strengthening transparency and trust in AI systems, RAIDO offers a practical foundation for organisations seeking to adopt more sustainable and reliable AI solutions. The pilot results also highlight the potential for the developed technologies to be adapted to other application domains beyond those explored within the project.



From Research Results to Market Uptake

Beyond the development and validation of its technologies, RAIDO has placed particular emphasis on the long-term exploitation of its research results. During the Final Industry Showcase, Eight Bells presented the project's approach to bringing these results closer to the market, introducing five Key Exploitable Results (KERs) that address different needs across the AI lifecycle. With Technology Readiness Levels (TRLs) ranging from 5 to 6, these results provide a foundation for further development, integration and adoption by AI developers, businesses, research organisations and other potential users.

The five Key Exploitable Results cover complementary aspects of AI development and deployment. The Data Enrichment and Distillation Software (KER1) addresses data scarcity and model efficiency, while the energy-efficient training toolbox (KER2) helps reduce computational and energy requirements. The RAIDO Green-AI Orchestrator (KER3) brings together different AI tools and optimises their workflows. To strengthen trust and transparency, verifiable AI model cards using blockchain technology (KER4) support the traceability of AI models, while a human-in-the-loop explainable AI interface (KER5) enables users to better understand AI-generated results and remain involved in the optimisation process.

To support the long-term sustainability of its results, RAIDO established a structured Intellectual Property Rights (IPR) and Exploitation Framework throughout the project's lifecycle. With the support of the Horizon Results Booster services, the consortium developed go-to-market strategies for its Key Exploitable Results, identifying target customers, potential business models and opportunities for integration into existing AI solutions. These efforts aim to facilitate the uptake of RAIDO's technologies beyond the project's completion, creating opportunities for further collaboration with industry,

research organisations and potential end-users.



Industry Engagement and Knowledge Exchange

The Final Industry Showcase also provided opportunities for direct interaction between the RAIDO consortium and participants through a dedicated poster and demonstration session. Attendees had the opportunity to explore the project's applications, discuss the pilot results directly with the teams and learn more about the potential use of RAIDO's technologies in their own organisations.



The morning programme also featured a panel discussion entitled *"What should responsible AI deliver by 2030?"*, moderated by Gohar Sargsyan from TCS, bringing together perspectives from industry, academia, UNESCO and the investment community. The session addressed key aspects of responsible AI, including energy efficiency, fairness and accountability, followed by a presentation by Paolo Ribeca from MPAI on the role of standardisation in supporting interoperable, trustworthy and reusable AI

technologies.



UCL UNESCO Chair on AI Forum

The afternoon programme continued with the UCL UNESCO Chair on AI Forum, connecting RAIDO's research objectives with the broader role of AI in sustainable development. Opened by Professor John Shawe-Taylor and María Pérez-Ortiz, the session highlighted the importance of responsible AI in addressing societal and environmental challenges. A presentation by Leona Isabelle Verdadero from UNESCO explored the partnership between UCL and UNESCO, with particular emphasis on Green AI and the role of academic research in informing international initiatives and policy development.

The afternoon sessions explored a range of topics, including AI bias, trustworthy climate information, critical infrastructure, education and urban resilience. The programme concluded with a panel discussion entitled *"AI for People and Planet – What Should Responsible AI Deliver by 2030?"*, bringing together different research perspectives on how AI can contribute to sustainable development and address real-world societal needs.



As RAIDO approaches the completion of its three-year journey, the Final Industry Showcase marked an important milestone in bringing European research closer to industry and society. Building on its validated technologies, pilot results and exploitation strategy, the project lays the groundwork for the wider adoption of trustworthy, energy-efficient and sustainable AI solutions, contributing to a future where AI works for both people and planet.

About RAIDO

RAIDO is a Horizon Europe-funded research and innovation project focused on advancing Trustworthy and Green AI through an integrated framework covering data management, model optimisation, explainability, and sustainable AI deployment. The project develops innovative methods for automated data curation and enrichment, including Digital Twins and diffusion models, alongside data- and compute-efficient AI techniques such as few-/zero-shot learning, continual learning, and model distillation. RAIDO further integrates explainability, blockchain-enabled traceability, reinforcement learning, and AI orchestration mechanisms to support transparent, accountable, and energy-efficient AI systems across cloud-edge environments. The platform is being validated through four real-world demonstrators in application domains including smart grids, smart farming, healthcare, and robotics.

Learn more and be part of the RAIDO community!

https://raido-project.eu/

https://www.linkedin.com/company/raido-project/

https://x.com/RAIDO_project

https://www.youtube.com/@RAIDOproject

RAIDO project has received funding under the European Union's Horizon Europe research and innovation programme under Grant Agreement No 101135800. UK participants in Horizon Europe Project RAIDO are supported by UKRI.

Co-funded by
the European Union