humAIne

Hybrid Human-AI Decision Support for Enhanced Human Empowerment in Dynamic Situations

Human-in-the-Loop AI for Trustworthy Decision Support: Insights from the HumAIne Pilots

ENFIELD Workshop on Human-Centric AI
Bucharest, Romania (Hybrid) | 09/09/2025

Vassilis Voulgarakis (OKYS)

OKYS
CLOUD COMPUTING APPLICATIONS

Funded by the European Union

Funded by the European Union. 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.

The HumAIne project in a nutshell

ℹ️ **A 3-year Horizon Europe initiative** dedicated to pioneering the collaboration between humans and AI-systems for decision-making in **dynamic, unstructured environments**.

🎯 The project aims to empower AI solution integrators through **the HumAIne Operating System (OS)** a platform that:

- ✓ integrates cutting-edge technologies into **a single system to surpass both standalone AI and Human capabilities** and
- ✓ transforms the landscape of decision-making and problem-solving **across five 'Smart Industry' sectors.**

🤝 **The team:**

Funded by the European Union. 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.

The project's scope

MOTIVATIONS

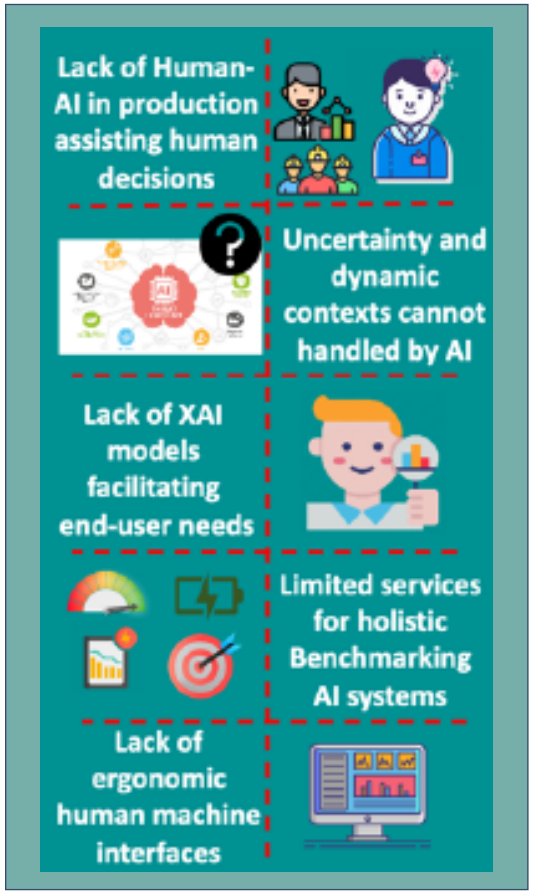

SOLUTION

HumAIne Innovative Objectives & Ambition

Unified platform for Human-centred AI decision-making
From single-aspect Human-AI to holistic & dynamic human-machine collaboration

Trustworthy & transparent and interactions between humans & AI
From local to global aware XAI systems considering biases & semantics

Active Learning for effective Human-in-the-Loop systems
Dynamic human enrolment towards reduced uncertainty in AI predictions

Swarm Learning for collective human intelligence
Increased AI accuracy and human empowerment with collective decision-making

Neuro-Symbolic AI for truly human and energy efficient AI
Efficient completion of complex tasks combining AI with semantic reasoning

Human-Machine Interfaces
Enhanced visualisations of AI outcomes interaction between human and machines, and shared knowledge .

Benchmarking Suite for Human-AI systems
Multi-aspect benchmarking of Human-AI applications

Novel Human-AI pilots and Use Cases
Validation of HumAIne KERs in real-life use cases from various high impact domains

EXPECTED IMPACTS

1. Supporting a human-centered and ethical development of digital industrial technologies
2. Contribution to the development of trusted and transparent AI-based systems
3. Humans can take full advantage of the benefits from AI systems
4. Enable collective intelligence for humans and AI systems using Swarm Learning
5. Accurate and trusted predictions of HumAIne systems leveraging domain knowledge based on Neuro-Symbolic AI

Funded by the European Union

Funded by the European Union. 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.

Keynote's fundamental thematic

⚠️ AI may risk replacing human judgment if not designed carefully

🤔 The key issues revolve around:

- **Trust by human factors**
- **Transparency of the AI-systems/models**
- **Accountability of the AI-systems/models**

AI should **augment** rather than replacing human decisions!

✅ So…**why human-in-the-loop?**

✓ Affords human factors to take control when AI confidence is low

✓ Ensure contestability and interpretability of AI-systems/models' outcomes

✓ The experts' feedback adapts systems/models dynamically

Human intervention is **a design feature**, not an afterthought!

Funded by the European Union

Funded by the European Union. 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.

The HumAIne OS and why it matters

✓ Active Learning

Humans intervene in AI processes when faced with increased uncertainty

✓ Neuro-Symbolic Learning

Deep Learning is combined with semantics and rules to complete complex tasks with high accuracy and minimal training

✓ Swarm Learning

AI models are trained in a decentralized and collaborative way across nodes without sharing raw data (privacy preserved)

✓ Explainable AI (XAI)

XAI models deliver holistic predictions and allow for users' feedback

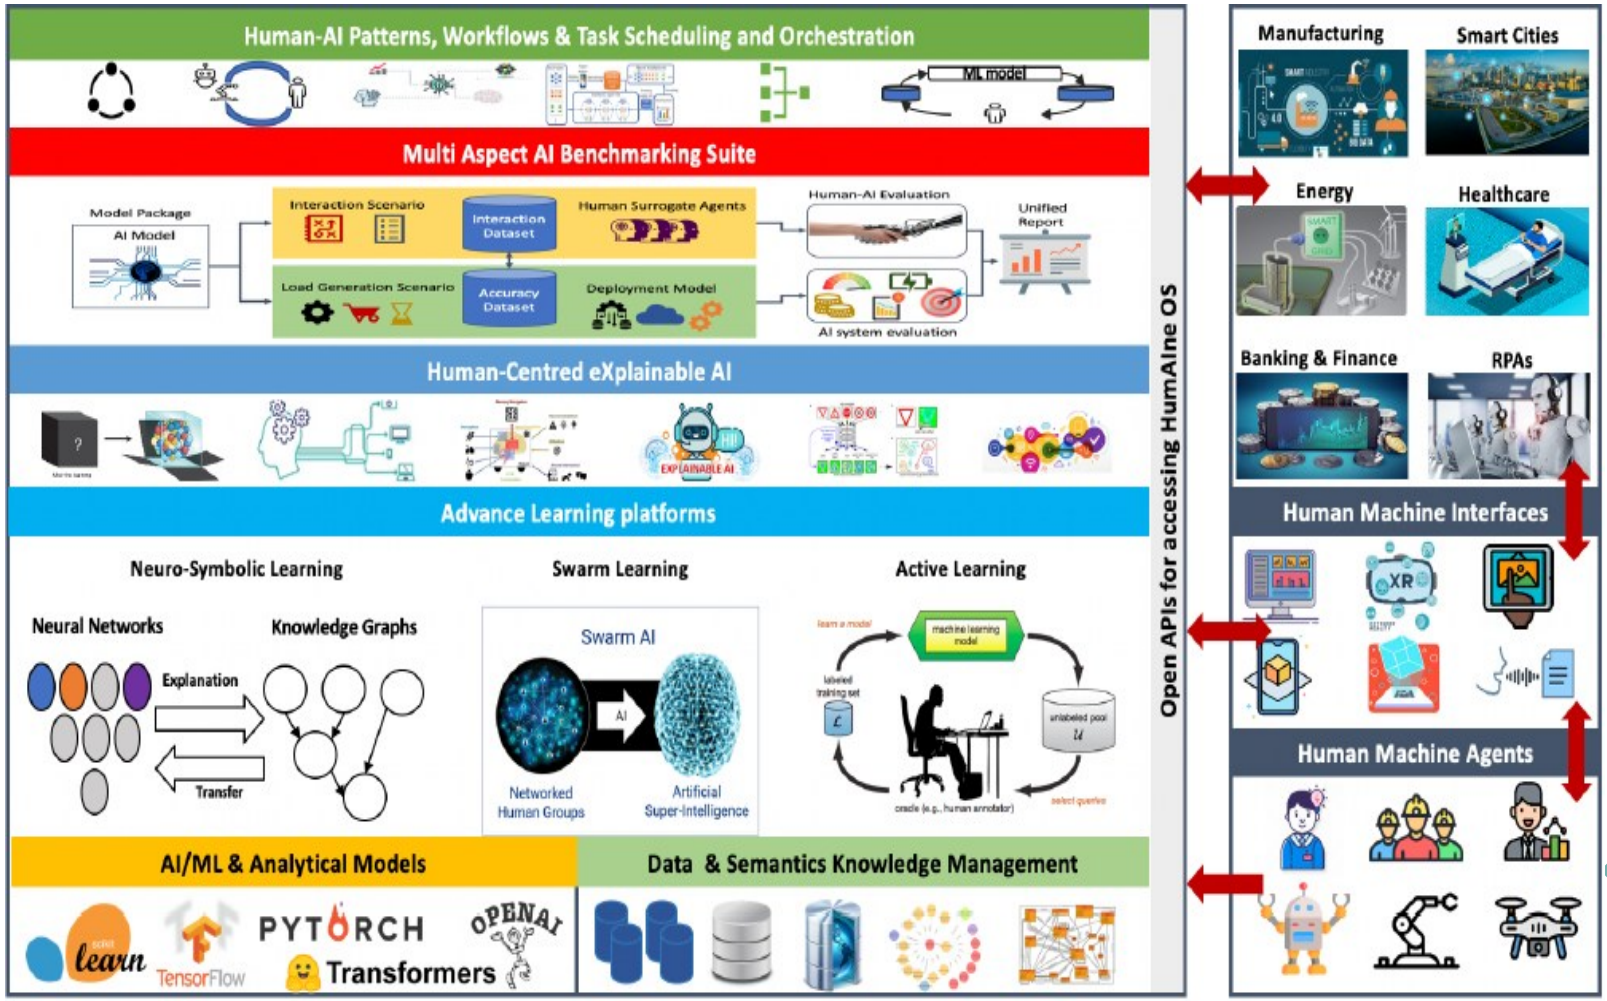

Funded by the European Union

Funded by the European Union. 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.

The HumAIne OS Architecture

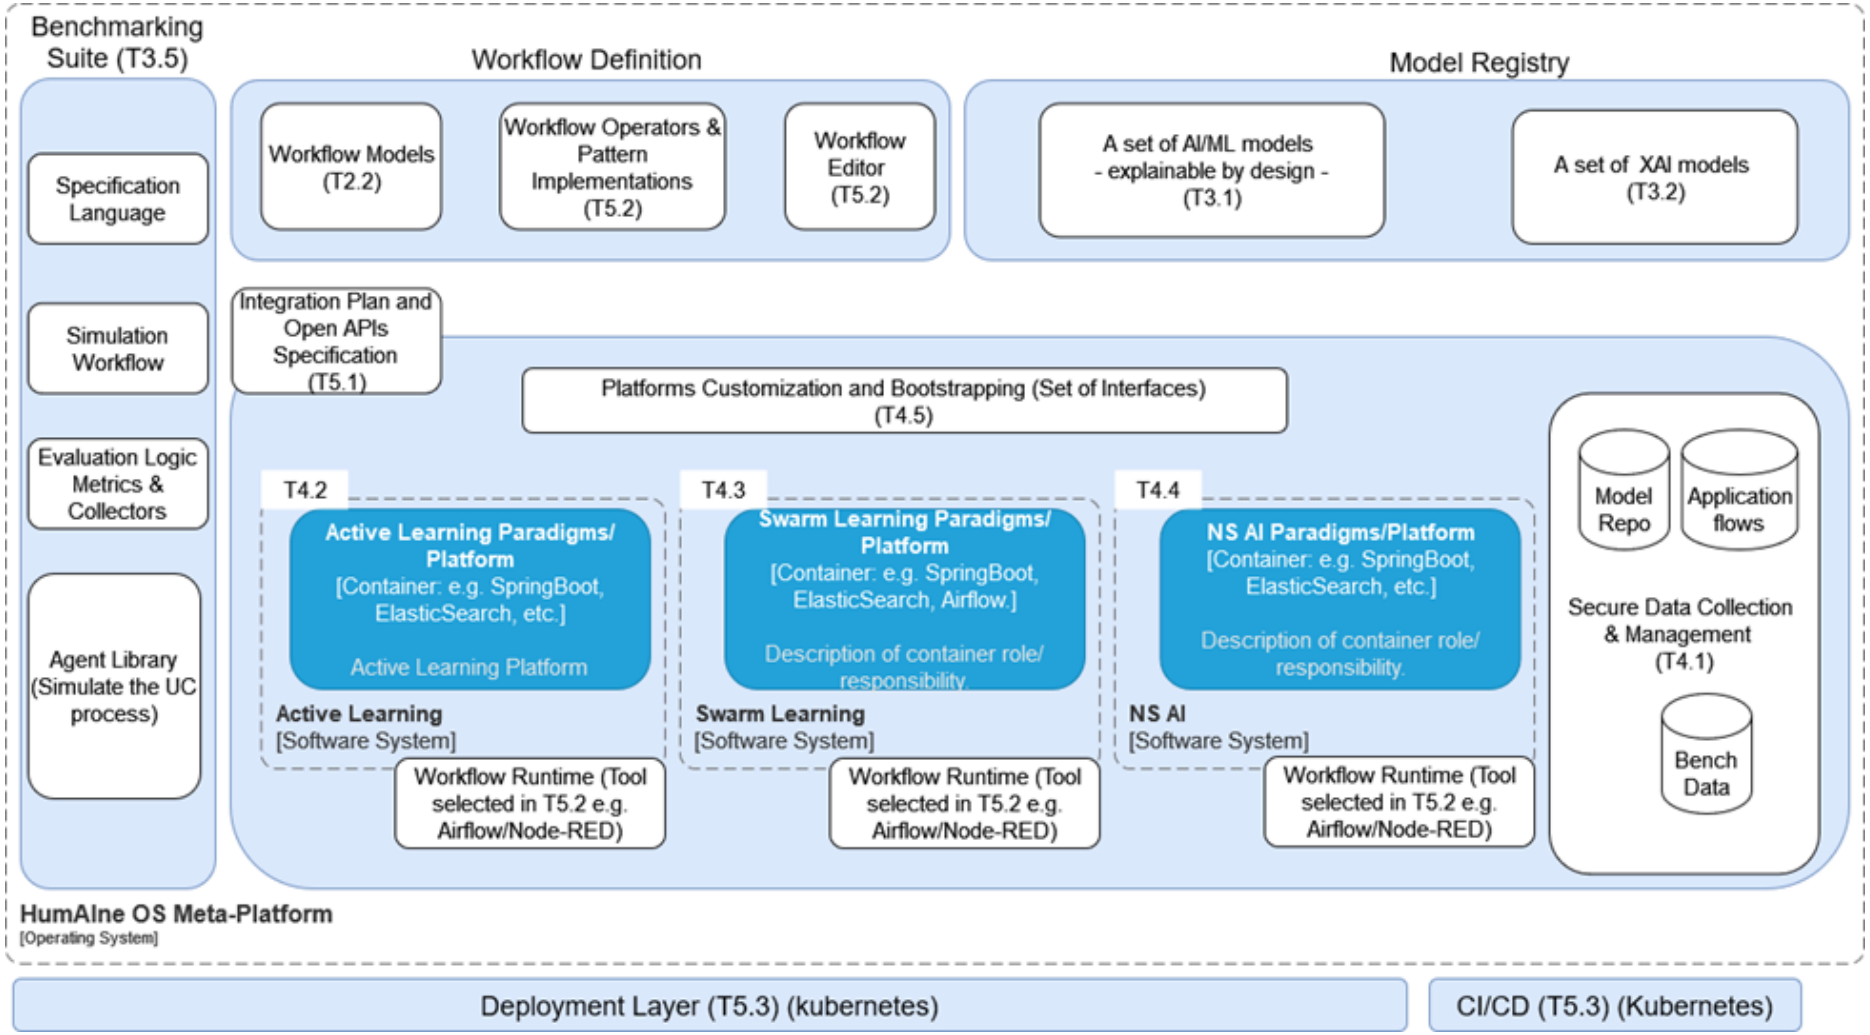

Funded by the European Union

Funded by the European Union. 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.

Transforming decision-making across Smart Industries

humAIne

Smart Manufacturing

Scheduling in a Machine-as-a-Service scenario

Smart Energy

AI-assisted decision support tools for system operators

Smart Cities

Robotic automation of processing parking permits

Smart Ticketing

Automating technical ticket resolution

Smart Healthcare (Diabetes)

Living at home without exacerbations & re-hospitalisations

Smart Healthcare (Oncology)

Managing patient workflow with AI-assisted tools

Funded by the European Union

Funded by the European Union. 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.

Smart Manufacturing Pilot

The context

Challenge:

- Complex production lines with modular tasks need optimal scheduling. Traditional scheduling methods are time-consuming and lack flexibility.
- Data ownership and privacy are key challenges in industrial settings.

AI Paradigm: Swarm Learning

- Distributed machine learning that keeps data local, ensuring privacy. AI model training is decentralized, with each machine as an agent.

Human Role:

- **Mixed Reality (MR) UI** ensures transparency and allows human operators to adjust goals and oversee AI predictions in real-time.

The Human-in-the-Loop in action!

Swarm agents learn from local data without transferring sensitive information.

MR Interface allows real-time human intervention, adjusting scheduling objectives as needed.

Efficiency improvement: Human oversight ensures accurate alignment with operational goals while maintaining privacy.

Privacy-preserving AI provides valuable insights without compromising data security.

Funded by the European Union

Funded by the European Union. 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.

Smart Energy Pilot

The context

Challenge:

- Power grid operations are increasingly complex due to renewable energy integration. Operators need real-time insights to make decisions based on accurate, interpretable predictions.

AI Paradigm: Active Learning + XAI

- AI helps forecast grid states and identify risks.
- Digital Twin simulates scenarios for decision-making, while XAI explains AI predictions.

Human Role:

- Operators interact with the AI outputs, leveraging AR/VR HMI for immersive decision support.
- Human feedback ensures that predictions and actions are interpreted correctly in dynamic grid conditions.

The Human-in-the-Loop feature

Digital Twin simulates the grid and helps operators visualize potential scenarios.

XAI explains forecasting errors and provides confidence levels for predictions.

AR/VR HMI allows operators to interact with simulations and AI insights, guiding manual interventions as needed.

Funded by the European Union

Funded by the European Union. 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.

Smart Cities Pilot

The context

Challenge:

- Manual processing of resident permit applications is labor-intensive and prone to errors. There's a need for automation, but with human supervision to ensure accuracy.

AI Paradigm: Active Learning

- The system continuously learns from operator feedback to improve accuracy. AI assists in identifying fields in PDFs, automating the data validation process.

Human Role:

- Council operators validate AI output and annotate training data, making corrections where necessary.
- Expert feedback directly influences AI learning cycles and improves model precision over time.

The Human-in-the-Loop feature

Active Learning enables operators to intervene when AI uncertainty is detected.

Operators annotate fields in PDFs to correct and fine-tune AI predictions.

Automation: AI automates routine validation checks, saving time and reducing errors.

Outcome: Processing time for a permit reduced from **20 minutes** to **10 seconds**, with humans ensuring ultimate control and accuracy.

Funded by the European Union

Funded by the European Union. 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.

Smart Ticketing Pilot

The context

Challenge:

- IT ticket systems face challenges in efficiently categorizing and resolving complex requests
- Human expertise is essential for effective dispatching, but AI can automate routine tasks to improve efficiency.

AI Paradigm: Active Learning + XAI

- AI continuously learns from human feedback to categorize tickets and suggest resolutions. XAI ensures that IT experts understand AI decisions and can intervene when necessary.

Human Role:

- IT experts validate AI predictions, help categorize new issues, and correct any errors improving AI performance over time.

The Human-in-the-Loop feature

Active Learning: AI improves by learning from IT expert feedback on ticket categorization.

XAI: Explains AI decisions, enabling experts to validate and correct outputs.

Outcome: AI models adapt continuously through expert input, reducing manual categorization time. Trust is built via transparent, explainable AI.

Funded by the European Union

Funded by the European Union. 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.

Smart Healthcare (Diabetes) Pilot

The context

Challenge:

- Personalized recommendations for diabetes management require real-time adaptation to patient needs. Traditional models rely on expert rules or data.

AI Paradigm: Neuro-Symbolic Learning

- Combines rule-based expert knowledge with data-driven AI learning for robust decision support. NSL helps predict outcomes based on past patient data and expert guidelines.

Human Role:

- Clinicians input expert rules and feedback on model predictions, ensuring personalized and safe care. The system adapts dynamically to clinician feedback, enhancing its accuracy over time.

The Human-in-the-Loop feature

NSL combines data-driven learning with expert feedback to continuously improve diabetes management strategies.

Clinicians can **intervene to adjust the system's decisions** based on individual patient circumstances.

Outcome: **Early results show +10% accuracy improvement using NSL** over baseline models, thanks to clinician input.

Funded by the European Union

Funded by the European Union. 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.

Smart Healthcare (Oncology) Pilot

The context

Challenge:

- Structured reporting in oncology is essential for diagnosis, but annotation is slow and inconsistent. AI models need to be accurate and interpretable, especially when dealing with health data.

AI Paradigm: Active Learning + XAI

- AL reduces manual labeling workload, and XAI improves transparency by explaining model decisions. XAI can flag uncertainty in AI-generated reports, so that radiologists maintain control.

Human Role:

- Radiologists provide input by annotating images and reports, so that AI models learn from correct examples. XAI highlights uncertainty, prompting

The Human-in-the-Loop feature

AL significantly reduces annotation load for clinicians by automating parts of the labeling process.

XAI ensures clinicians can verify AI outputs and identify potential issues with model predictions.

Outcome: Faster report generation, with clinicians still overseeing critical aspects to ensure accuracy and relevance.

Funded by the European Union

Funded by the European Union. 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.

Key Principles for Effective Human-in-the-Loop 1/2

Feedback Aggregation & Consolidation:
Collects human feedback across multiple interactions, grouping it to guide system learning.
✓ Example: In **Smart Cities**, operators annotate data in batches, improving AI accuracy over time.

Expert Consensus Consolidation:
When multiple experts provide feedback, the system consolidates diverse opinions to ensure a unified decision.
✓ Example: In **Healthcare – Oncology**, radiologists validate AI predictions, with any disagreement triggering further review and retraining.

Low-Confidence Alerts:
The system notifies humans when AI confidence drops below a threshold, prompting intervention.
✓ Example: In **Smart Manufacturing**, AI alerts operators when schedules are uncertain, allowing them to adjust goals in real-time.

Funded by the European Union

Funded by the European Union. 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.

Key Principles for Effective Human-in-the-Loop 2/2

Continuous Learning & Adaptation:
AI models adapt and improve through ongoing expert interaction, ensuring they evolve with real-world data and feedback.

- ✓ Example: In **Smart Energy**, operators adjust forecasts via AR/VR interfaces, guiding AI toward better predictions.

Transparency & Explainability:
AI decisions are explained clearly to users, fostering trust and allowing for meaningful human intervention.

- ✓ Example: In **Smart Ticketing**, XAI explains categorization choices, ensuring IT experts can trust and correct the model's decisions.

Funded by the European Union

Funded by the European Union. 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.

Key takeaways & lessons learnt

🙏 **Trust requires explainability and usable interfaces**

AI models are trusted more when users understand how they make decisions and can interact with them in an intuitive way.

👌 **Adoption depends on empowering, not replacing experts**

AI-systems/models perform best when human input helps guide learning, especially in situations of uncertainty or complex decision-making.

Human feedback must be integrated quickly and efficiently into AI models to continuously improve performance without overwhelming users.

🔁 **Scalability and Flexibility are essential for successful integration**

Human-in-the-loop systems must be adaptable and scalable across various industries to meet the diverse needs of different sectors

Funded by the European Union

Funded by the European Union. 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.

humAIne

Hybrid Human-AI Decision Support for Enhanced Human Empowerment in Dynamic Situations

Thank you for your attention!

humaine-horizon.eu

info@humaine-horizon.eu

Funded by the European Union

Funded by the European Union. 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.