AI Ethics from Principles to Practice

What ten Industrial AI Deployments tell us about making AI Ethics work

DATE
01 July 2026

AUTHORS
Horizon Europe Project AIOLIA (Operationalising AI Ethics for Learning and Practice: A Global Approach)
Grant Agreement N° 101187937

Key message

Ethics operationalisation is essential, beneficial, and possible for organisations – but it requires structural support. The AI Act creates obligations without providing tools for their implementation, while ALTAI is too abstract for industry to implement AI ethics in real life. Organisations, particularly smaller ones, are left to develop processes from scratch, often without adequate resources, guidance, or coordination infrastructure to support meaningful implementation of ethics in their AI designs and deployments. This brief sets out six concrete recommendations to address these issues.

1. Problem and policy context

Context

Across Europe, organisations are deploying AI systems that directly influence how people think, decide, and behave. The EU AI Act sets requirements for ethical AI – but turning those requirements into practice is far harder than the law acknowledges. The Assessment List for Trustworthy AI (ALTAI), published by the EU's AI High-Level Expert Group in 2020, remains the most widely used self-assessment tool for organisations deploying AI. Its seven requirements (covering human agency and oversight, technical robustness, privacy, transparency, fairness, societal well-being, and accountability) provide a broadly applicable framework.

What AIOLIA adds

AIOLIA's research shows that ALTAI is too abstract for AI developers and has structural gaps in practice. The consequence is that it has only limited value for organisations in implementing AI ethics when they develop or deploy AI.

AIOLIA created a unique, bottom-up ethics process for organisations to ensure ethics is deployed in practice. It was created together with industry and translates ALTAI's high-level ethics principles into concrete, measurable technical and organisational measures within a deployment context. The work further resulted in a portfolio of 175 practical measures, 12 contextualised ethics principles, 37 components, and 6 process recommendations – grounded in real-world experience, not theoretical ideals.

175	12	6
Practical measures	**Ethics principles**	**Process recommendations**
Technical & organisational, across ten use cases	*Bottom-up, contextualised from real deployments*	*On how to operationalise ethics effectively*

Ten use cases, one shared challenge

AIOLIA created a framework for the operationalisation of AI ethics in cooperation with industry partners in ten use cases: six European use cases (UC1-UC6) and four non-European use cases (UC7-10) spanning professional and personal AI deployment contexts. Five use cases operate in professional settings (UC1–UC4, UC7) and five in personal contexts (UC5, UC6, UC8-10). The deployment context is significant: AIOLIA demonstrates that professional deployment contexts emphasise accountability, transparency, and decision traceability, while personal deployment contexts foreground user autonomy, safety, and long-term well-being. This distinction has direct implications for how ethics frameworks, including the AI Act and ALTAI, are applied.

Ethics operationalisation is effective risk management

A consistent insight across all ten use cases is that organisations approach AI ethics primarily through the perspective of risk management and benefits – not abstract moral values. Organisations that embed ethics can gain concrete advances: reduced legal exposure, stronger stakeholder trust, enhanced reputational legitimacy, and new business opportunities as ethics-aware providers. Policy-makers should communicate the positive business case, not only the compliance burden of ethics.

High-level principles are necessary but insufficient

AIOLIA's bottom-up co-creation process confirmed that well-established ethics frameworks – including the ALTAI Assessment List and OECD AI Principles – resonate with practitioners. However, abstract principles alone are unable to guide the concrete implementation of ethics in their AI designs or deployments. Practitioners need the translation of these principles into context-specific actions. Current frameworks do not bring ethics to this operational level.

Ethics principles interact, create tensions, and require managing trade-offs

One of the most practically significant findings of AIOLIA is that ethics principles do not function independently. Across all ten use cases, partners identified interdependencies and tensions amongst principles. Transparency requirements, for instance, can conflict with privacy protections, while safety monitoring requires data collection that clashes with user anonymity. These tensions cannot be resolved by addressing abstract ethics principles in isolation; they require ongoing institutional mechanisms to negotiate how to balance competing ethics expectations and manage trade-offs.

AI-driven change in professional behaviour vs individual cognition

AIOLIA reveals a structural difference in how ethics principles play out depending on deployment context. In professional settings (healthcare, safety engineering, HR, security), the core concerns are accountability, traceability, and preventing over-reliance on AI in

high-stakes decisions. In personal settings (AI companions, personal assistants, deepfake therapy), the main focus is on user autonomy, psychological safety, and long-term well-being. Current AI governance frameworks, including ALTAI, do not adequately reflect this distinction – and the AI Act's high-risk classification system does not fully capture the risks in personal AI use.

2. Policy recommendations

1	**Develop implementation-driven guidance to complement the AI Act's principles**	The AI Act (recital 27) establishes what organisations must achieve in AI ethics but not how. AIOLIA demonstrates that the way obligations like 'ensure human oversight' or 'provide explainability' translate into practice varies enormously by context, sector, and AI system type. The European AI Office and national supervisory authorities should develop sector-specific, implementation-driven guidance (drawing on practitioner co-creation processes like AIOLIA's) that shows organisations concretely how to operationalise each requirement, not merely what the requirement says.
2	**Update ALTAI to reflect the professional/personal context distinction, temporal risks and practical handling of competing ethics requirements**	AIOLIA's bottom-up process reveals important gaps in ALTAI that limit its usefulness for practitioners. ALTAI does not adequately distinguish between professional and personal AI use contexts or account for the long-term and cumulative effects of AI on human cognition and behaviour (such as deskilling, dependency, and gradual manipulation). It further lacks guidance on the context-specificity of ethics implementations across the AI lifecycle does not adequately account for tensions when implementing competing ethics principles. The European Commission should provide accompanying guidance to ALTAI (treating AIOLIA's D3.1 + D3.2 as direct inputs) to address these gaps. The AI Office should consider developing ALTAI variants for specific high-risk AI systems under Annex III of the AI Act.
3	**Resource and expand regulatory sandboxes to support ethics operationalisation for SMEs**	Small and medium-sized organisations often lack the capacity to build comprehensive ethics governance processes from scratch. While Article 57 of the AI Act requires member states to establish regulatory sandboxes, these are primarily framed around legal compliance testing. Policymakers should expand the scope and resourcing of sandboxes, and leverage Digital Innovation Hubs, to offer structured support for ethics operationalisation, including access to multidisciplinary expertise, co-creation facilitation, and independent review mechanisms that are affordable to smaller actors.
4	**Mandate and standardise ethics process requirements, not only outcome requirements**	The AI Act focuses primarily on the outputs and properties of AI systems (accuracy, explainability, non-discrimination). AIOLIA demonstrates that the effective implementation of ethics depends strongly on the process through which ethics is operationalised: who is involved, how trade-offs are managed, how feedback is collected, and how findings are used to update systems and procedures. Policymakers should introduce process-level requirements for high-risk AI – mandating co-creation with affected communities, multidisciplinary governance committees, feedback loops, and documented ethics decision-making – alongside existing output requirements.
5	**Recognise and fund the co-creation model as a standard practice for ethics operationalisation**	AIOLIA's co-creation process, pairing academic ethics and AI expertise with industrial partners through structured operationalisation cycles, produced 175 concrete measures that neither party could have developed alone. AIOLIA has condensed them into a 6-page *Measures Portfolio*. This Portfolio generates ethics guidance that is simultaneously principled and practically grounded. Policymakers should recognise co-creation partnerships as a legitimate and fundable approach to AI ethics implementation, and create EU-level infrastructure (funding, coordination, methodology toolkits) to make this model accessible beyond research projects.
6	**Invest in longitudinal research on AI's impacts on human cognition and behaviour**	AIOLIA identifies a critical knowledge gap: long-term effects of AI systems on human cognitive capacities, decision-making patterns, social skills, and well-being are poorly understood. Without this evidence base, organisations and policymakers are making precautionary judgements without solid empirical basis. The EU should fund structured longitudinal research programmes, enabling independent researchers to access operational AI systems and usage data, with robust data protection safeguards, to build the evidence based needed for informed, proportionate regulation.

3. Policy Implications

Implementation timeline

Immediate

Recommendations 1 and 2 (implementation guidance and ALTAI update) are immediately actionable. They can be initiated by the European Commission in 2026.

Midterm

Recommendations 3 and 5 (SME support and co-creation model) require medium-term investment (2026–2028).

Legislative

Recommendation 4 (process requirements) involves legislative and regulatory development and is a medium-term priority.

Continuous

Recommendation 6 (longitudinal research) should begin immediately and be funded continuously across the legislative cycle.

The compliance framing is both a risk and an opportunity

AIOLIA shows that industrial partners are supportive of AI ethics, especially when it is framed as approach to risk management and business benefit rather than moral obligation. Policymakers can harness this orientation by clearly communicating the legal, reputational, and commercial risks of non-compliance, while simultaneously emphasising the business benefits of implementing ethics compliance. If the AI Act is primarily framed as a compliance burden, this motivation will be lost.

Voluntary measures cannot substitute for structural support

Across all ten use cases, partners implemented ethics measures that exceeded legal requirements. But they did so as individual organisations, without shared infrastructure, standardised tools, or access to common expertise. The result is often a patchwork of siloed innovation and knowledge. Structural policy support (e.g., guidance, sandboxes, co-creation infrastructure, funded research) is necessary to bring ethics operationalisation into standard practice.

Potential barriers

- **Industry heterogeneity:** a single implementation approach cannot serve all sectors; guidance must be context-sensitive, requiring sustained co-development with diverse industry actors.
- **Regulatory capacity and bandwidth:** developing sector-specific implementation guidance at scale requires significant expertise and resources for stakeholder consultation.
- **Research access:** independent longitudinal research on the impact of AI on users requires AI providers to grant data access, which raises liability and IP concerns that must be addressed through carefully designed legal frameworks.
- **Fast moving developments:** as AI systems evolve rapidly, implementation guidance and ALTAI updates risk to be become outdated. A regular review cycle must be built in from the outset.

Conclusion

Ethics is a continuous reflective and anticipatory practice, embedded in organisational culture, governance structures, and industrial, professional contexts. The EU AI Act and ALTAI have created important foundations for AI ethics. What is needed is an operational translation of the regulatory framework to help organisations meet their obligations: practical guidance, shared tools, co-creation support, skills infrastructure, and a sustained evidence base on AI's impacts. AIOLIA provides a tested model and a ready-to-use portfolio of measures. The six recommendations in this policy brief show how policymakers can make them work in practice.

Horizon Europe AIOLIA project (www.aiolia.eu)

Coordinator CEA

Lead Partner CENTRIC

Funded by the European Union under Horizon Europe Grant Agreement N° 101187937. Views and opinions expressed are those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Executive Agency (REA). Neither the European Union nor the granting authority can be held responsible for them.

This brief draws on

Bayerl, P.S., Lawlor, E., Maris, M.T., Miorandi, D., Pekšys, G., Bjelica, M., Stojšin, K., Anastasova, M., Smith, O., Henestrosa, A., Yamshchikov, I., Bak, M.A.R., & Akhgar, B. (2026). Operational Ethics Guidelines on Use Cases Related to Human Behaviour and Cognition. AIOLIA Deliverable D3.1.

Cossette-Lefebvre, H., He, G., Huang, L., Ide, K., Katirai, A., Kim, K., Kim, J., Kim, G., Kishimoto, A., Maclure, J., Mikami, K., Suh, J., & Zhao, Y. (2026). Operational context-sensitive guidelines in Canada, China, Japan, and South Korea. AIOLIA Deliverable D3.2.