AI ETHICS AND INTEGRITY
INTERNATIONAL ASSOCIATION

JULY 2026

AI HORIZON
JOURNAL

#4
EDITION

CLOSING THE COMPLIANCE GAP: A NEW TOOL FOR EU AI ACT READINESS

A NEW FREE SELF-ASSESSMENT PLATFORM FROM AIEI TO MEASURE YOUR ORGANISATION'S EU AI ACT COMPLIANCE READINESS

UPDATED AIEI PRINCIPLES: STRENGTHENING RESPONSIBLE AI PRACTICES

THE REVISED AIEI AI PRINCIPLES AND THE NEW DECLARATIVE AI CERTIFICATE

PREPLY: SMART LEARNING, RESPONSIBLE TECH

HOW PREPLY BALANCES HUMAN OVERSIGHT WITH ADVANCED AI

0 23991 18471 9

FEATURING PREPLY | CALMA | UNICORN FACTORY | MINISTRY OF DIGITAL TRANSFORMATION OF UKRAINE

CONTENTS

JULY 2026 / ISSUE NO.04

INTRODUCTION

MESSAG

ASSOCIATION UPDATES

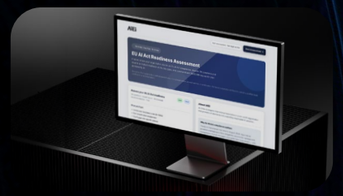

ALLIANCE UPDATES

COVER STORY

INDUSTRY INSIGHTS

WELCOME MESSAGE

There is a particular tension running through this edition — between the sophistication of the conversation around AI governance and the stubborn simplicity of what most organisations still lack: a clear picture of where they actually stand.

The EU AI Act is no longer a framework in progress. It is in force. Its prohibitions are already binding. Its obligations for general-purpose AI models are live. And in less than one month, the full weight of high-risk AI requirements becomes enforceable — with penalties that are not symbolic. Yet the gap between regulatory expectation and organisational readiness remains, for most, uncomfortably wide. Not because the rules are unclear. But because self-knowledge is harder than it sounds. That gap is precisely why we are launching the AIEI EU AI Act Readiness Assessment — a free, self-managed tool that returns something genuinely useful: a readiness score, a gap analysis, and a structured action plan. You cannot close a gap you have not mapped. The rest of this edition makes the same argument from different directions.

Borys Pratsiuk of Preply shows what it looks like when an organisation has done the harder work — building AI that genuinely serves its users rather than replacing them, maintaining balanced datasets, drawing clear lines between what the algorithm decides and what humans must. It is a case study in the kind of integrity that governance frameworks try to mandate but cannot manufacture.

From Kyiv, Oleksandr Bornyakov, Ukraine's Acting Minister of Digital Transformation, offers something rarer still: a government that is not just complying with AI regulation but actively trying to shape it. The transition from Digital State to Agentic State — with a sovereign language model, a national AI assistant already used by nearly half a million citizens, and a regulatory sandbox with nineteen companies — is a proof of concept worth studying, regardless of where you sit.

Ruslan Drozdov of Calma names something most business leaders will recognise but few have said plainly: that the compulsive monitoring of AI tool launches is not competitive intelligence. It is a tax. The FOMO Tax, as he calls it, costs not just time but leadership attention — and the operators genuinely ahead are not the ones testing the most tools, but the ones who chose fewer, went deeper, and built measurable processes around what they kept.

Gil Azevedo from Unicorn Factory Lisboa rounds out the practitioner voices with a view from one of Europe's most active startup ecosystems — one where ethical AI is not a constraint on competitiveness but, increasingly, a precondition for it.

Our Legal Committee members, meanwhile, report from the frontlines of regulatory implementation across four jurisdictions: Romania's missed deadlines and underfunded oversight machinery; Portugal's structural blind spot in criminal procedure, where defendants have no mechanism to scrutinise AI-generated evidence used against them; Singapore's world-first framework for agentic AI systems; and Scotland's ambition to govern AI for economic growth across an eight-layer national strategy. These are not abstract policy debates. They are the operating conditions your organisation is navigating right now.

Finally, we are publishing the updated AIEI Principles — ten commitments covering the full AI lifecycle, developed with more than twelve experts across nine countries — alongside the AIEI Declarative AI Certificate for organisations that want to signal their alignment with responsible AI practices in a market where that signal is beginning to matter.

The thread connecting all of it is the same: accountability has arrived as a present condition, not a future obligation. The question is whether you engage with it on your own terms — with adequate preparation and honest self-assessment — or wait until the gap between where you are and where you need to be becomes someone else's problem to manage for you.

Start with the assessment. Everything else follows from knowing the truth of where you stand.

AI ETHICS AND INTEGRITY INTERNATIONAL ASSOCIATION

EU AI ACT
READINESS ASSESSMENT

FREE, ONLINE TOOL

To read more about the platform or start assessment — visit **www.assessment.ai-ei.org**

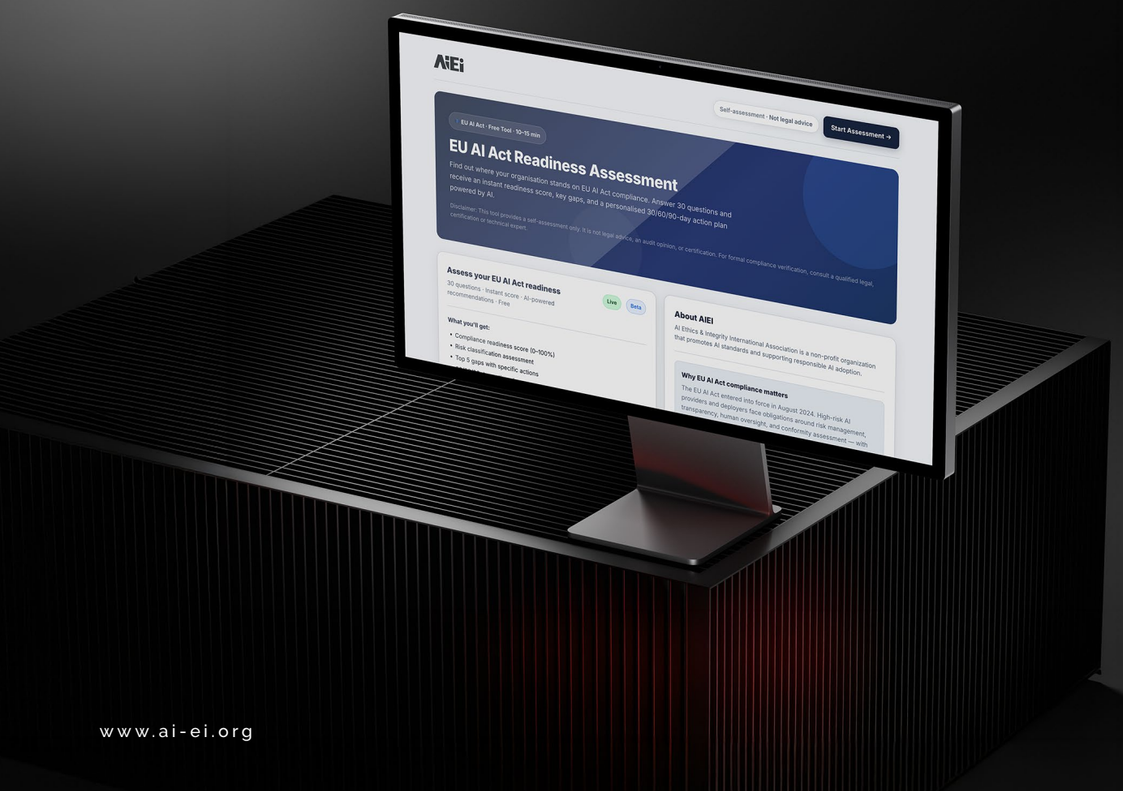

CLOSING THE COMPLIANCE GAP: A NEW TOOL FOR EU AI ACT READINESS

A new self-assessment platform from AI Ethics and Integrity International Association offers organisations a fast, free, and actionable way to measure their EU AI Act compliance readiness — no registration, no legal jargon, no barriers.

The EU AI Act is no longer a distant regulatory horizon. Having entered into force in August 2024, its provisions are rolling out in stages — with prohibitions on unacceptable-risk AI systems already in effect since February 2025, general-purpose AI model obligations active from August 2025, and full high-risk AI obligations kicking in by August 2026. For many organisations, the compliance clock is ticking, but the path forward remains unclear.

UNDERSTANDING THE EU AI ACT

The EU AI Act is the world's first comprehensive legal framework for artificial intelligence, adopted by the European Union to ensure that AI systems placed on the European market are safe, transparent, and respect fundamental rights. Rather than regulating AI as a single category, the Act takes a risk-based approach — classifying AI systems into four tiers depending on the potential harm they pose.

At the top sit systems deemed to carry unacceptable risk, such as social scoring by governments or real-time biometric surveillance in public spaces — these are outright prohibited. Below that, high-risk AI systems — covering areas such as hiring, credit scoring, education, critical infrastructure, law enforcement, and medical devices — face the most demanding obligations. A broader set of systems must meet transparency requirements, ensuring users know when they are interacting with AI.

The Act applies not only to companies based in the EU, but to any organisation worldwide whose AI systems are used within the European Union — making its reach genuinely global.

KEY COMPLIANCE TIMELINE

The Act's obligations are being phased in over a two-year period:

- August 2024 — The Act entered into force
- February 2025 — Prohibitions on unacceptable-risk AI systems became applicable
- August 2025 — Rules for general-purpose AI (GPAI) models, including foundation models, took effect
- August 2026 — Full obligations for high-risk AI systems become enforceable

With the final and most consequential deadline now less than one month away, organisations that have not yet begun their compliance work are running out of runway.

WHAT THE PLATFORM DOES

The tool guides users through 30 structured questions across six critical compliance domains: AI System Classification, Risk Management, Data Governance, Transparency, Human Oversight, and Conformity Assessment. In approximately 10 to 15 minutes, participants receive an instant readiness score on a 0–100% scale, along with a personalised report that includes:

- A risk classification assessment for their AI systems
- The top five compliance gaps with specific, actionable remediation steps
- A 30/60/90-day action plan generated by AI
- An evidence and documentation checklist
- Guidance on how AIEI can further support their compliance journey
- The entire process requires no registration and is provided completely free of charge.

WHO IT IS FOR

The platform is designed for any organisation that develops, deploys, or uses AI systems within the scope of the EU AI Act — from startups and scale-ups building AI products, to large enterprises integrating third-party AI tools, to public sector bodies navigating high-risk use cases.

WHY IT MATTERS NOW

The EU AI Act introduces substantial obligations for providers and deployers of high-risk AI systems: mandatory risk management processes, robust data governance practices, transparency obligations toward end users, and human oversight mechanisms — with penalties for non-compliance that can reach €35 million or 7% of global annual turnover, depending on the violation.

Yet many organisations, particularly those outside the largest enterprises, lack the internal legal or technical expertise to rapidly assess where they stand. The AIEI platform democratises access to that first critical step: self-awareness.

As AIEI states, the tool is not a substitute for formal legal advice, an audit, or certification — and it makes this clear upfront. What it does provide is something equally valuable at this stage of the regulatory rollout: a structured, credible, and immediately actionable baseline assessment that organisations can act on today.

The assessment is live and available at
www. assessment.ai-ei.org

AI ETHICS AND INTEGRITY INTERNATIONAL ASSOCIATION

LEGAL COMMITTEE INSIGHTS

MEMBER UPDATES & ANALYSIS

Our Legal Committee members are continuously monitoring AI regulatory developments across jurisdictions — tracking how governments, courts, and institutions are translating AI policy ambitions into real legal frameworks.

Vitor Neves identifies a structural gap in Portugal's AI regulatory picture: while supervisory authorities have been designated and sectoral coordination is underway, the country's criminal justice system contains no rules for handling AI-generated evidence — leaving defendants without meaningful recourse when algorithmic outputs are used against them. He argues that market surveillance and procedural fairness are being treated as separate problems, when in practice they are two halves of the same compliance challenge.

Kateryna Kernoz examines Scotland's newly published AI Strategy 2026–2031, unpacking how the government plans to balance economic growth with responsible adoption — and what the strategy's eight-layer framework actually means for businesses and SMEs still hesitant to engage with AI. She pays particular attention to the gap between the strategy's ambitions and current adoption rates, and what practical steps might close it.

Alexandru-Mihai Alexandrescu looks at Romania's uneven path toward EU AI Act compliance — a country whose national strategy speaks confidently of digital transformation while its supervisory infrastructure remains critically underfunded and months behind on key deadlines. His analysis weighs whether the government's pragmatic reliance on existing institutions can hold up once full enforcement begins in August.

Finally, Looi Teck Kheong turns to Singapore, where a deliberately flexible, sector-by-sector approach to AI governance is taking shape — most notably through the world's first framework specifically designed for agentic AI systems, unveiled at Davos in January 2026. His analysis explores why Singapore's model of voluntary frameworks paired with targeted legal updates may offer a compelling alternative to comprehensive

PORTUGAL'S AI REGULATORY FRAMEWORK: A CRIMINAL JUSTICE BLIND SPOT

Vitor Neves
President of Associação Jurídica do Porto (AJP), AIEI Legal Committee Member

As the EU AI Act moves toward full application in August 2026, most attention has focused on market surveillance, penalty frameworks, and national authority designations. Portugal is no exception, yet its situation reveals a blind spot that goes far deeper than administrative compliance. While the country has designated ANACOM as its market surveillance authority and national point of contact and coordinates oversight across fourteen sectoral bodies, its criminal justice system remains entirely unprepared for the procedural challenges that high-risk AI systems create. Portuguese criminal procedure contains no rules on algorithmic evidence, no right to explanation when AI outputs are used to build suspicion or proof, and no mechanism for the defence to scrutinise the technical foundations of systems used by investigating authorities. The result is a structural imbalance. In June 2026 the Prosecutor General issued an AI ethics charter — prosecutors may not use AI to predict recidivism or recommend coercive measures — but it only constrains how the prosecution uses the technology. It gives the accused nothing. When a facial-recognition match reaches the file, the defence has no right to an explanation and no way to scrutinise the system behind it. Italy has legislated on courtroom AI; Spain and Germany have implementation laws under way. Portugal's only answer has come from the prosecution, about the prosecution.

To read full article, please scan this QR codeor visit this page.

SCOTLAND'S AI STRATEGY 2026–2031: GOVERNING AI FOR GROWTH

Kateryna Kernoz
Attorney-at-Law (Ukraine), LL.M. in Technology and AI Law (UK), AIEI Associate Member

In March 2026, the Scottish Government published its AI Strategy 2026–2031, presenting artificial intelligence as a driver of economic growth, productivity and public service transformation. Structured around an eight-layer "AI Stack", the strategy outlines actions across business adoption, infrastructure, data and regulation while maintaining a commitment to responsible and trustworthy AI. It proposes measures ranging from AI adoption support for SMEs and skills initiatives to investment in infrastructure and closer alignment with international regulatory principles. With 61.9% of surveyed Scottish SMEs reportedly not using AI technologies, the strategy also identifies business adoption and AI literacy as key priorities for the years ahead.

To read full article, please scan this QR code or visit this page.

ROMANIA AND AI REGULATION: BETWEEN EUROPEAN ALIGNMENT AND NATIONAL READINESS

Alexandru-Mihai Alexandrescu
Associate Lawyer at Buju, Stanciu & Associates, AIEI Associate Member

The EU AI Act has been in force since August 2024, and its obligations are already binding on companies operating in Romania with or without a fully functional national supervisory framework. Yet Romania only took its first formal implementation step in March 2026, missing the authority designation deadline by over six months. The proposed model assigns oversight to existing institutions, with ANCOM as the lead authority coordinating a network of sectoral regulators. It is a pragmatic choice, but the Government's own memorandum flags a sobering constraint: a shortage of the financial and human resources needed to supervise complex AI systems effectively. Meanwhile, Romania's National AI Strategy 2024–2027 speaks confidently of human-centric values, regulatory sandboxes and digital transformation. No dedicated national AI law exists, legislative proposals remain at an early stage, full applicability is set for August 2026.

To read full article, please scan this QR code or visit this page.

AI REGULATION IN SINGAPORE: A 2026 UPDATE

Looi Teck Kheong
Global AI Ambassador and President of the Singapore Chapter at the Global Council for Responsible AI; Advocate & Solicitor of the Supreme Court of Singapore. AIEI Legal Committee Member

While the EU pursues comprehensive regulation and other jurisdictions scramble to designate supervisory authorities, Singapore governs AI through voluntary frameworks, sectoral guidelines, and targeted updates to existing law. The approach is deliberate: flexibility for industry, accountability through existing legal channels, and trust as a competitive advantage. The most notable development of 2026 is the Model AI Governance Framework for Agentic AI, unveiled at the World Economic Forum in Davos in January by Singapore's Minister for Digital Development and Information and developed by IMDA — the first framework of its kind globally to address AI systems capable of autonomous reasoning, planning, and independent action. Alongside it, Singapore announced plans for a Digital Infrastructure Act imposing binding requirements on data centres and cloud providers, and — building on its Generative AI Evaluation Sandbox, launched in 2023 in partnership with Google, Microsoft, Anthropic, and Amazon Web Services — continues to develop common evaluation standards for AI.

To read full article, please scan this QR code or visit this page.

AI, TRUST AND THE PLACE WHERE THE HUMAN STAYS

AI Ethics & Integrity International Association
at the Data with Purpose Summit

www.ai-ei.org

A KEYNOTE RECAP FROM THE DATA WITH PURPOSE SUMMIT, PORTUGAL

At the Data with Purpose Summit in Oeiras, Portugal, AIEI President Sergiy Barbashyn argued that the hardest problem in artificial intelligence is no longer technology, but trust — and set out what responsible AI looks like in practice.

Here is a version of the AI story that most people in the industry tell themselves: the technology is difficult, the engineering is hard, and once those problems are solved, adoption will follow naturally. Sergiy Barbashyn, President of the AI Ethics and Integrity International Association, opened his keynote at the Data with Purpose Summit in Oeiras with a quiet dismantling of that assumption. The capability question, he argued, is largely settled. The harder problem — the one that does not yield to more compute or faster deployment — is trust.

The gap is not subtle. Stanford's AI Index puts it in stark relief: roughly three-quarters of AI experts expect the technology to improve how people work. Among the general public, that figure falls to around one in four. That distance is not a communication problem to be solved with better marketing. It is a structural challenge — and it is now the thing worth competing on.

He grounded the argument in something deliberately unglamorous: ordinary moments where AI touches ordinary lives. A loan application declined without explanation. An insurance claim rejected by a system whose logic no one can surface. A job candidate screened out before a human ever sees their name. The technology works; the harder question is whether people are willing to live with what it does.

THE LIMITS OF THE RULEBOOK

Regulation, in Sergiy's framing, is a necessary but insufficient response. Rules will always trail technology. The EU AI Act is the most serious attempt yet to set limits that hold — and it matters, not least because its penalties are real and its reach is global. But no regulation can substitute for a habit of judgment.

"A system should be judged by how it's used, not by the word 'AI,'" Barbashyn said. The examples he chose were instructive in their ordinariness: a workplace wellbeing app that begins reading employees' emotional states; a child-safety feature that quietly shapes what minors see; a training dataset that resurfaces years later as a copyright dispute no one anticipated. None of these is a dramatic failure. Each is a case where a reasonable-sounding early decision compounded into a problem that proved difficult to undo.

RESPONSIBILITY DOES NOT DISAPPEAR — IT MOVES

There is a tempting but mistaken idea that introducing AI into a process transfers responsibility from humans to machines. Sergiy rejected it plainly. "AI doesn't remove human responsibility. It relocates it." When a model denies a loan, a human designed the model, chose the training data, and decided which decisions the system would make without review. The accountability does not evaporate with automation. It shifts upstream.

This is the argument that connects ethics to durability. Organisations that treat responsible AI as a constraint on what they can build are misreading the situation. The practical moves are within reach of any organisation: keep a real human in control of decisions that affect real people, collect only the data you actually need, and be able to say plainly where AI is used and who is accountable for it. These are not heroic commitments. They are the baseline the moment requires.

WHAT AIEI IS FOR

Through its AI Principles and Declarative Certificate, the Association offers organisations a framework for making commitment visible and accountable — before a regulator asks. Through the AI Horizon Journal and annual Conference, it surfaces real-world cases that rarely make it into formal policy documents. And through the EU AI Act Readiness Assessment — a free, self-managed tool returning an instant compliance score, a gap analysis, and a structured action plan — it gives organisations the one thing good intentions cannot substitute for: an honest picture of where they actually stand.

Sergiy closed on a note that had the quality of a provocation disguised as an observation. "Technology is not neutral. It carries the values of whoever builds it." Every design decision is a value decision. Every choice about what data to use, what outcome to optimise for, what human to keep in the loop — encodes something about what the organisation believes. The question is not whether those values are present. It is whether they are examined.

That, in the end, is what the trust gap is about. Not a failure of technology. A failure of attention.

Sergiy Barbashyn is President of the AI Ethics and Integrity International Association. The Data with Purpose Summit took place in Oeiras, Portugal. For more information, visit ai-ei.org.

UPDATED AIEI PRINCIPLES: STRENGTHENING RESPONSIBLE AI PRACTICES

Our AI Principles have been updated. This latest version reflects the Association's ongoing commitment to translating complex ethical and regulatory concepts into clear, actionable guidance that organizations can genuinely apply.

The updated AIEI AI Principles reflect the Association's continued effort to ensure that its framework for responsible artificial intelligence remains practical, relevant and aligned with evolving international approaches.

The revised Principles provide organizations with a structured and consistent approach to the responsible use, development, and implementation of AI systems. They are designed to support organizations across the entire AI lifecycle, from design and procurement to deployment, monitoring and retirement, while maintaining an appropriate balance between innovation, human rights and ethical responsibility.

WHY DO THESE PRINCIPLES MATTER?

AI integration into business processes is no longer a question of the future, it is the reality of today. With that comes growing responsibility: for data quality, algorithmic transparency and the real-world consequences of automated decisions on people's lives.

The AIEI Principles provide a voluntary framework that helps organisations clearly articulate their approach to AI – from design and procurement through to monitoring and retirement. They do not replace applicable law, but offer a practical reference point for responsible AI use across the full system lifecycle, maintaining a balance between innovation, human rights, and ethical responsibility.

CORE PRINCIPLES FOR RESPONSIBLE AI

The framework covers the entire AI lifecycle and is structured around ten interconnected principles:

1. **Respect for Human Rights and Ethics.** AI systems should respect fundamental rights, freedoms and human dignity, ensuring compliance with legal and ethical standards. Organizations should assess potential impacts on individuals, especially in sensitive contexts.
2. **Human Oversight.** AI should support, not replace, human decision-making. Meaningful human control must be ensured, particularly in high-impact or high-risk scenarios.
3. **Fairness and Non-Discrimination.** AI systems should avoid bias and ensure equitable outcomes. Organizations should actively identify and mitigate risks of discrimination in data and algorithms.
4. **Privacy and Data Protection.** Personal data is collected only where there is a valid legal basis and a legitimate purpose. Individuals' data rights must be respected throughout the AI lifecycle.
5. **Transparency (Explainability).** AI use should be clearly disclosed, and key decisions should be explainable in understandable terms. Stakeholders should be informed about how AI systems operate.
6. **Safety, Security, Robustness and Reliability.** AI systems must function reliably and securely, with safeguards against failures, misuse, or attacks. Continuous monitoring and improvement are required.
7. **Accountability and Governance.** Clear lines of responsibility must be established, including for third-party AI tools obtained from vendors, contractors, or external platforms.
8. **Societal Benefit and Responsibility.** AI should be used in ways that benefit society and avoid harm. Organizations should consider broader societal impacts beyond immediate business value.
9. **Environmental Sustainability (Green AI).** AI development should take into account environmental impact, promoting energy efficiency and sustainable practices where possible.
10. **Awareness, Education and Workforce Adaptation.** Organizations should promote AI literacy and support workforce adaptation. Employees must understand how to use AI responsibly and effectively.

AIEI DECLARATIVE AI CERTIFICATE

Members of the Association who confirm their support for the Principles may obtain the AIEI Declarative AI Certificate (AI Verified badge). The certificate:

- Reflects an organization's commitment to responsible AI practices
- Supports transparency and credibility in the market
- Serves as a visible signal of alignment with ethical and governance standards

The initiative is voluntary and does not impose regulatory obligations. Instead, it is intended to encourage responsible practices and continuous improvement.

The updated version of the Principles was developed with the contribution of more than 12 experts from 9 countries, led by the AIEI Legal Committee including Amisha Mittal, Kateryna Kernoz, Hysmir Idrizi, Jae-Seong Lee, Taras Lytovchenko and to all contributors involved.

THE ROLE OF THE STATE IN THE DEVELOPMENT OF AI

An Interview with Oleksandr Bornyakov, Acting Minister of Digital Transformation of Ukraine

As artificial intelligence reshapes modern governance, Ukraine has emerged as a compelling case study in state-led AI adoption — forged not in peacetime stability, but amid the pressures of an ongoing war. Oleksandr Bornyakov, Acting Minister of Digital Transformation of Ukraine, outlines his ministry's vision: the transition from a "Digital State" to an "Agentic State," in which AI is embedded in the very logic of how government operates.

From a sovereign national language model built on Ukrainian cultural memory, to the citizen-facing platform Diia.AI, Ukraine is positioning itself not merely as a recipient of Western digital norms, but as an active architect of Europe's emerging AI regulatory landscape — offering real-world proof of concepts that balance bold transformation with democratic accountability.

1. HOW WOULD YOU DESCRIBE THE KEY ROLE OF THE MINISTRY OF DIGITAL TRANSFORMATION OF UKRAINE IN FORMING THE AI ECOSYSTEM OF UKRAINE TODAY? WHAT AREAS OF WORK ARE THE MOST PRIORITY?

The Ministry of Digital Transformation of Ukraine plays a role as the architect of national technological sovereignty. Ukraine is currently transitioning from a Digital State to an Agentic State. To realise this goal, the Ministry launched the WINWIN AI Center of Excellence — the first in Europe to comprehensively integrate AI into the state. The center has already launched three national projects and is developing 18 state AI products, eight of which are already in use.

Today, we are focusing on four strategic areas. The first is creating AI Factory with NVIDIA to provide sovereign computing power. The second is developing a national language model based on Google's Gemma architecture, trained on Ukrainian registries, archives, scientific articles and media, preserving Ukrainian cultural context and remaining free from Russian narratives. The third is the AI Development Strategy until 2030, covering public administration, education, defense and medicine while also strengthening infrastructure, preparing state data and expanding AI expertise. The fourth is integrating AI agents into citizens' daily lives through Diia.AI, the world's first national AI assistant for government services, already used by more than 450 thousand citizens.

The Ministry is also working on AI regulation using a bottom-up approach: first preparing companies for future requirements, and only then adopting legislation. On May 15, 2025, Ukraine signed the world's first international agreement in the field of AI – the Council of Europe Framework Convention.

2. WHAT IS THE CURRENT STAGE OF THE AI DEVELOPMENT ROADMAP? WHAT ARE THE MAIN PLANS AND PRIORITIES OF THE MINISTRY FOR THE NEXT 3–5 YEARS?

Our ambition for the next 3–5 years is to become one of the top three global leaders in AI development by 2030. This is being realised through four vectors.

The first is building "Smart AI-Sovereignty": together with Kyivstar we are developing the national language model Syaivo and launching AI Factory with NVIDIA. The second is scaling the Agentic State: the state must become a proactive assistant, including with voice functionality through a partnership with ElevenLabs. The third is DefenseTech and technological self-sufficiency. The fourth is synchronization with the EU and Big Tech: full compliance with the European AI Act, but not simply copying norms — rather influencing international legislation through our own real-world cases.

3. WHAT IS THE MINISTRY'S ROLE IN CREATING A REGULATORY FRAMEWORK FOR AI? HOW CAN IT STIMULATE INNOVATION WHILE FORMING A FOUNDATION FOR HARMONIZATION WITH EUROPEAN STANDARDS?

We launched the Sandbox – a special space for safely testing AI and blockchain innovations without the risk of fines, with 19 companies already joined. The key challenge is harmonization with GDPR. Our approach is "smart AI sovereignty", where we clearly distinguish between open data and information requiring total protection.

But there is one fundamental point for us. In the process of digital European integration, Ukraine is not just a student blindly copying others' norms. But an active architect of Europe's future single digital space. We have the ambition to influence international legislation — including the European AI Act — through our own real-world proofs of concept. We want to show the world in practice how to structure rules that guarantee data security without killing bold AI transformation.

4. WHAT APPROACHES TO INTERNATIONAL COOPERATION DO YOU USE TO UNIFY LEGISLATION AND ESTABLISH THE PRINCIPLES OF RESPONSIBLE AI?

Our main approach is active partnership and leadership through cases. Ukraine has signed the Council of Europe Framework Convention on Artificial Intelligence, Human Rights, Democracy and the Rule of Law. We are implementing the HUDERIA methodology as a tool for assessing the impact of AI systems on human rights, democracy and the rule of law. Thanks to our negotiations, OpenAI and PayPal officially entered the Ukrainian market and we launched joint projects with Microsoft, NVIDIA, ElevenLabs and Palantir.

5. TELL US MORE ABOUT THE AI SANDBOX CONCEPT. WHAT REAL BUSINESS CHALLENGES DOES THE PROJECT HELP SOLVE AND CAN IT HELP THE STATE BETTER UNDERSTAND BUSINESS NEEDS?

The Sandbox is a joint project of the Ministry of Digital Economy and the Ukrainian Startup Fund – a platform where companies can test their solutions, receive expert advice and verify compliance with national and international legislation such as the AI Act or GDPR. For businesses, it reduces the risk of errors, saves time and resources, and builds trust with partners and investors. For the state, it provides hands-on experience with specific products and cases, enabling the creation of a clear and effective legal framework.

6. WHAT ARE THE MAIN CHALLENGES (TECHNICAL OR LEGISLATIVE) THAT YOU SEE IN IMPLEMENTING AI AT THE STATE LEVEL?

The first technical challenge is data quality. Ukraine has around 350 state registers, and some still require significant improvement – that is why we are actively modernizing our backend. The second is security and sovereignty: we are creating AI Factory and the national language model Syaivo to process data exclusively within the country. The legislative challenge is striking a balance between rules and innovation. We have chosen a bottom-up approach: instead of imposing strict bans, we first give businesses tools to self-regulate through the AI White Paper, guidelines and regulatory sandboxes.

7. DOES THE STATE PLAN TO CREATE SPECIAL CONDITIONS (ON THE BASIS OF DIIA.CITY OR SEPARATELY) SPECIFICALLY FOR R&D CENTERS IN THE FIELD OF ARTIFICIAL INTELLIGENCE?

We do not plan to create a separate regime for AI – Diia.City already ideally suits this purpose, uniting over 4,200 residents including more than 200 companies in R&D and over 200 in AI. In parallel, we are creating physical infrastructure: the Brave1 cluster for defense tech startups and RoboLab for hardware and robotics developers. For AI developers, the most important thing right now is access to computing power – that is why we are working on the AI Factory project.

To read full interview with Oleksandr Bornyakov, please scan this QR code or visit this page.

Ministry of Digital Transformation of Ukraine

Oleksandr Bornyakov
Acting Minister
of Digital Transformation of Ukraine

Diia City

INTERVIEW

SMART LEARNING, RESPONSIBLE TECH

How Preply Balances Human oversight with Advanced AI

Guest

BORYS PRATSIUK

Engineering Manager at Preply

Preply

Borys's background includes 10+ years of experience in AI and management. Starting as an Android developer, he later held key leadership roles, including Head of R&D at Ciklum and CTO at Scalarr. Now, Borys focuses on leading engineers, driving engineering success and bridging technical gaps to ensure seamless execution across teams.

Interviewer

SERGIY BARBASHYN

AIEI President

AIEI AI ETHICS AND INTEGRITY INTERNATIONAL ASSOCIATION

Accomplished legal expert with over 13 years of experience in IT, AI, and IP law. Sergiy is active member of numerous professional associations, including UNBA, AIPPI, EYBA, INTA, and Board. His portfolio includes serving as President of the European Young Bar Association, Head of Board Legal, Deputy Chair of the IP Committee at Ukrainian National Bar Association, and Chair of the UNBA NextGen.

Sergiy Barbashyn, President of the AI Ethics and Integrity International Association (AIEI), sat down recently with Borys Pratsiuk for a conversation that cuts to the center of a question most executive teams are now wrestling with: how do you scale AI without losing sight of the people it's meant to serve?

The timing is not incidental. Across nearly every sector, AI has moved from a pilot-stage experiment to a core operational layer, embedded in decision-making, product design, and customer experience alike. Leaders are being asked to make calls on AI governance, data stewardship, and algorithmic accountability with the same rigor they'd apply to financial or regulatory risk. Few playbooks exist. Fewer still have been tested at scale.

At the AIEI, we've spent considerable time trying to give that conversation some structure. Our framework rests on two pillars: Ethics, which we define as a commitment to human-centric values, and Integrity, which we define as the uncompromising quality of technical execution behind those values. The two aren't separable. A company can publish all the right principles and still fail its users if the underlying systems are sloppy, opaque, or under-tested. Equally, a technically flawless system built without regard for the people affected by it isn't a success. It's a liability waiting to surface.

That framework is what brought Sergiy to this conversation with Borys. Preply, as a company operating at the intersection of education and AI-driven personalization, sits in an interesting position: its product touches real learning outcomes for real people, at a scale where small design decisions compound quickly. We wanted to understand what principles guide that work internally, how those principles hold up under commercial pressure, and where they intersect with the broader case the AIEI has been making about ethical AI deployment.

Borys was generous with his time and candid in his answers. In the tradeoffs, the internal debates, and the decisions that don't make it into a mission statement but shape the product anyway.

Borys, let's start with some highlights about Preply. How can we describe the Platform in a few sentences?

Preply is an international platform for private online lessons, founded in 2012, now having offices in Ukraine, the USA, Spain and the UK. The vision is to power language learning progress by connecting students with the best tutors for a truly personalized experience. Preply operates across 180 countries. Also, Preply provides 90+ languages and 163 subjects. There are 100,000 tutors teaching millions of lessons per month. And all numbers are growing continuously.

Thank you. We have checked that there are other players in your market as well, such as Duolingo, Italki and Babbel. May you clarify the key distinctive components that distinguish Preply from these market players?

Not all of these companies we consider as direct competitors, as their services significantly differ from what we do. Preply actually helps people speak and unlock language capabilities. This growth was driven by an ambitious and bold strategy that integrated AI tools with a human-first approach.

Key innovations include a lesson recap, a self-practice app for vocabulary and a new AI avatar, which will help students to practice daily conversations. And lastly, we prioritize the continuous development of our tutors. Our dedicated department helps tutors to level up their skills through formal certifications and quality approvals. By investing in our teachers' expertise, we aim to upgrade the quality of lessons and results of every student.

Moving closer to the AI topic, what are Preply's key AI strategic directions?

One of our main priorities is analyzing data from lessons to accelerate student learning. We treat this with the utmost responsibility, ensuring mandatory user approval. The data belongs to our users and we just save it in a safe way, analyze and use it to make students` learning with Preply faster.

The second direction focuses on internal ecosystem automation. By integrating AI into coding, communication and other core business processes, we make operations at Preply faster, smarter and more efficient.

Finally, our applied research group works on its own models and leverages data to implement cutting-edge solutions like unique algorithms for analyzing pronunciation mistakes.

Thank you, that is really impressive. Speaking of unique algorithms, our readers would like to know more about "matching algorithms" at Preply. How does the system connect students with the right tutors?

Preply's powerful matching algorithm analyzes multiple parameters to connect students with the best tutor for them. We evaluate how lessons are conducted, focusing on student outcomes and tracking regional patterns and learning habits. Also, we have noticed that if the tutor starts a lesson with a "good morning" or "how are you", that improves student retention.

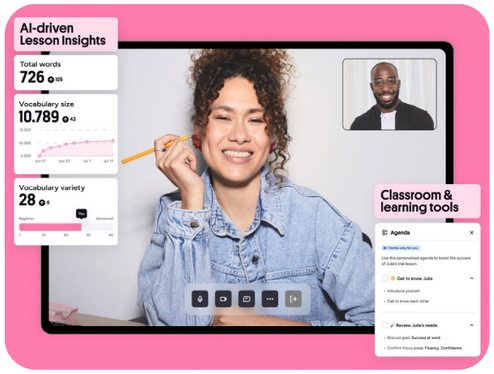

And if a person is looking for a gentle tutor with good manners, we can easily connect this intelligence to starting a lesson with “hello”, “good morning”, “how are you”. Our system is evolving constantly, ensuring that we respect and leverage each tutor’s unique background.

Also, each tutor at Preply has some skills in specific areas, like finance, law, IT, medicine and some people specifically search for such a tutor for different reasons. It could be for enrolling in university, improving or changing a career, etc. Preply allows you to book just a couple of lessons with a tutor from a certain industry and go deep in the topics of your interest.

We have a real case at Preply with a Japanese student who worked in a globalization company in Japan. Despite having a limited budget, she was interested in finding out how her colleagues from the US work in the globalization area. So, she found a tutor from Preply who works in a globalization company in the US and uses the sessions to dive into American modern practices. Eventually, she created a manual and following it was really useful in her work. That is not just about language transfer, it’s actually more about knowledge transfer that addresses specific professional or educational gaps. Just for reference, 90% of lessons now are still language learning, but we are expanding into new, different topics, such as public speaking, psychology and many others. We try to maintain diversity and encourage its development by welcoming professionals from many fields and from all over the world to join Preply as tutors.

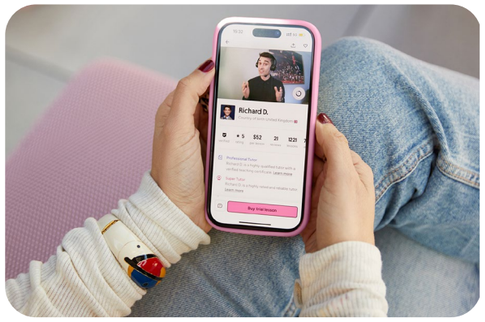

That’s a wonderful philosophy! You earlier mentioned the AI avatar and other AI tools at Preply. How do you divide tasks between those where human interaction is prioritized and those best optimized by AI? I would like to know what that synergy looks like at Preply.

We were surprised, but when we started doing different pilots, we found out that people choose Preply because it’s a non-AI avatar space. Students like to talk human to human and we make it possible by connecting people.

For lessons, we prioritize live human interaction, while AI is used for highlighting mistakes that might otherwise go unnoticed. At Preply, we ensure that AI acts as a copilot for human instruction, not competes or replaces humans. So we are exploring AI innovations for supplemental practice to help students improve some skills in between the lessons. For instance, with our AI avatar, you can easily practice ordering coffee or another real-life scenario. With the script, this tool guides the conversation and points you to continue dialogue in the same direction, making sure the student is not distracted by other topics. At Preply, we note that practising small talk with an AI Avatar in between the lessons shows really good results in language learning and in improving pronunciation for our students.

Another way we implement AI is transcription of lessons for the sake of history and as a data source. At Preply, AI can also be used to generate homework based on specific student mistakes. By reducing the time spent on manual exercise creation, we allow tutors to focus on reviewing and approving the most effective content. We believe that while AI provides a variety of tools, only the tutor truly understands what is best for the specific student’s progress and helps to motivate them throughout the learning process.

As we understand, AI can act as a copilot to help tutors and students with tasks like fixing grammar mistakes or summarizing. Nevertheless, how do you maintain a balanced dataset to ensure your AI tools are trained properly? Also, how do you avoid algorithmic bias?

We built our own grammar correction model that we use. When developing it, we focused on ensuring that the training and evaluation data reflected the diversity of our learner base. That includes learners across different proficiency levels, geographies, and demographic groups.

By maintaining a balanced dataset that represents equality and diversity, we prevent algorithmic bias and ensure high performance in production. If we build an unbalanced dataset, it will not work in production because of wrong results.

Regularly, we collaborate with OpenAI, which has been our trusted partner for a long time. In addition to our proprietary efforts, we leverage their API too. By doing this, we are able to bring a better learning experience faster to more and more people.

Though AI can be a nice assistant, it may also hallucinate. Is it applicable to your work in AI implementation?

So, in rare cases, hallucination might happen. However, we aim to avoid it by strictly controlling the main processes. For these needs, we have an internal evaluation framework that is based on manual review data. And the number of data points for evaluation is constantly growing. The special framework that we use is adopted by all teams at Preply, and that’s the key protection from possible hallucinations. However, we are still improving deeper ethical moments to ensure the result is sufficient and no hallucinations are observed.

From a privacy standpoint, that makes total sense. Looking at the learning process itself from another angle, do you track student progress at Preply and do you use AI for this purpose? We are especially curious about how you measure the quality of a student’s progress, not just the quantity of lessons taken.

We have a dedicated team that is working on progress tracking using metrics over time. Progress measurement is a really scientific task to do. We are still in progress and are constantly trying to implement new metrics.

By analyzing how students’ vocabulary changed from simple words to more complex ones, our AI provides a clear picture of their growth. Preply visualizes this journey through personalized charts, mapping the student’s unique learning path and the results they have achieved during their learning with us. But measuring the progress of other parts of language, like fluency, grammar, listening and reading abilities, is a more complex task, which we are actively working on.

Well and lastly, what is your point of view on AI regulation and governance in the modern world? What do you think should be on the table: regulation, soft framework or maybe guidance?

Preply takes our legal obligations seriously and treats GDPR as a baseline we build on, not just a box to tick. When it comes to regulation, a layered approach makes the most sense: binding rules where the risks to people are highest, and lighter frameworks and guidance where the technology is still evolving rapidly. We are transparent about what is recorded, and nothing is captured unless users actively choose to enable the AI feature themselves.

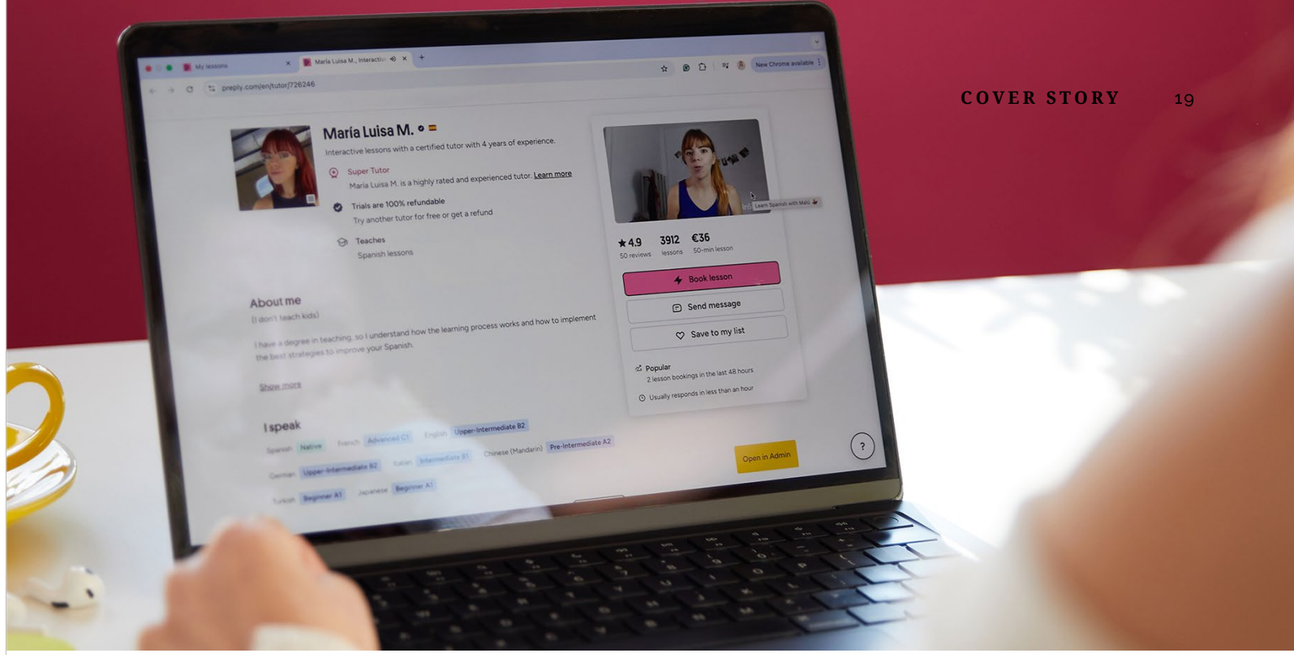

As AI becomes more capable, we’re increasingly faced with fundamental questions about where technology should assist humans and where humans should remain firmly in control. Our view is that AI is at its best when it helps people make better decisions, not when it replaces human judgment entirely.

We believe effective AI governance will require a combination of regulation, industry standards, and practical guidance. Regulation is important for protecting users and establishing accountability, but innovation also moves quickly, so companies need clear frameworks and best practices that can evolve alongside the technology.

Ultimately, the goal should be ensuring AI remains transparent, responsible, and human-centered as a tool that empowers people rather than one that makes important decisions on their behalf.

Thanks a lot for this insightful discussion. Let`s summarise with three notes from you on Preply’s approach with AI.

Firstly, at Preply, we are using data for good reasons, specifically to unlock the student`s learning velocity.

Secondly, we believe AI should enhance the learning experience in a responsible, inclusive, and human-centered way. For us, that means building AI-powered solutions that support learners and tutors across different backgrounds, goals, and learning needs, while keeping fairness, trust, and quality at the center of our approach.

And the third one, we continuously develop new AI-powered products, aiming to improve our product, learning processes and to mature the company`s technological capabilities. So AI is our big friend and we are using it for good reasons.

To read about Preply, please scan this QR code or visit preply.com

BUILDING IMPACT THROUGH RESPONSIBLE AI

Author
Gil Azevedo,
Executive Director
at Unicorn Factory Lisboa

ABOUT
UNICORN FACTORY LISBOA

Unicorn Factory Lisboa is a platform of programmes and innovation hubs designed to position Lisbon as a leading European capital of innovation.

Launched in 2022 by the Lisbon City Council, Unicorn Factory Lisboa develops a range of initiatives and programmes that support both national and international startups and scaleups.

With the goal of strengthening the city's entrepreneurial ecosystem and establishing Lisbon as a global innovation benchmark, Unicorn Factory Lisboa operates across five key areas: early-stage incubation for startups; growth-stage programmes for scaleups; soft-landing support for international startups and scaleups expanding into new markets; youth entrepreneurship; and a network of innovation hubs, including the Beato Innovation District, web3hub, gaminghub, greenhub, AIHub, healthub, and engineersHub.

The success of Unicorn Factory Lisboa played a key role in Lisbon being awarded the title of European Capital of Innovation 2023.

Europe is accelerating toward an AI-driven economy, and Portugal is no exception. Innovation has flourished increasingly across the country, creating impact across multiple sectors. Despite the common belief, companies are not the only ones benefiting from this technology. According to a BCG survey, the use of AI is growing among individuals, and nearly 80% of young people between the ages of 18 and 34 already use it.

Many companies, associations, and groups are looking on how to incorporate AI in their activities, making efforts to promote knowledge and organize new training initiatives. These initiatives are truly changing the way AI can impact our daily lives in a responsible manner. Concerns about its use are also growing legitimately, with 11% already fearing it may lead to a complete replacement of the workforce.

Unicorn Factory Lisboa is one of these key initiatives which is reshaping how innovation is built, governed, and scaled. Our mission is to accelerate the creation and growth of high-impact startups, reinforcing Lisbon's position as a global launchpad for technology companies. The results are clear, and we are proud of our path: Lisbon attracted more than 80 international tech companies in the recent years, including 17 unicorns, announcing over 17,000 jobs. The project has been paramount for the election of Lisbon as the European Capital of Innovation by the European Union in the end of 2023. Plus, we have also been featured as the highest-ranked portuguese organization in the Europe's Leading Start-up Hubs 2026 by Financial Times.

At the same time, we are aware of the challenges ahead, and the ethical use of AI is one of the most important focus areas. That is one of the main reasons that led to the creation of the AIhub, a dedicated physical space designed to support AI innovation and, at the same time, bridge the gap between AI ambition and responsible execution. The hub counts with the Center for Responsible AI, a consortium of AI startups, as a partner and it houses the Portugal AI Innovation Factory, led by Microsoft with Accenture, Avenade and Unicorn Factory Lisboa, a living lab where startups access not just workspace, but a complete ecosystem for ethical AI development.

From our experience working closely with numerous startups across different sectors in the AI space, these new entities face a new type of pressure: the need to innovate quickly, while maintaining responsibility. Bias in datasets, lack of transparency in model decisions, and the potential misuse of generative systems are no longer theoretical risks. For early-stage teams, these challenges are amplified by limited resources, increasing regulatory complexity, and rising societal expectations. Portugal, like Europe, still faces challenges with excess of bureaucracy and an excess of regulation that can limit the use of AI, leaving the lead in innovation to other geographies.

While it is key to simplify bureaucracy, responsible use of AI is more relevant than ever. Building compliant, secure, and transparent AI systems from day one strengthens user trust and accelerates responsible adoption and impact.

At Unicorn Factory Lisboa, we are developing a comprehensive ecosystem that help founders navigate AI regulation and scale responsibly. Through the AIhub, we keep exploring new initiatives with key partners, such as Google and Microsoft, providing cloud infrastructure, AI tooling, and enterprise-level technical mentorship. The Portugal AI Innovation Factory, already referred, offers structured programs including 1-2 day immersive workshops with industry experts, strategic roadmap sessions to align with EU AI Act requirements, and rapid prototyping sprints. The partnership with Center for Responsible AI is also fundamental to guide startups and new projects that want to use AI to create real impact in a responsible way, but do not know where to start.

Many startups in our AI portfolio have seen strong enterprise engagement, particularly where alignment with ethical AI framework requirements is considered important for partnerships, investor confidence, and market differentiation. For example, Noxus recently partnered with PwC, highlighting the importance of compliance in new projects. Other recent examples from our portfolio include AI-powered diagnostic systems being piloted in Portuguese healthcare institutions, such as Usawa Care or Rapidoc, as well as AI-driven educational platforms that personalize learning, like Growappy. We also have several other startups with strong AI trajectories, including Granter, which supports grant funding, and Brainr, an all-in-one AI-powered factory management platform for the food industry that boosts efficiency and productivity.

We recognize that each market has a different approach to AI, but we share a common objective that goes beyond borders: ensuring startups build solutions that deliver real impact. The ethical use of AI creates not only stronger products, but also stronger companies, with the potential to generate even more qualified jobs, improve quality of life, and contribute to more dynamic cities. AI is already accelerating medical discovery, improving diagnostics in healthcare, enabling more personalized education, and connecting communities in more effective ways. It is one of the most powerful engines of the modern economy, but with significant unknowns, which makes its ethical development essential.

Looking ahead, applied AI in essential sectors such as health, education, defense and public services will continue to expand, alongside growing demand for infrastructure such as compute access, data spaces, and trusted European AI frameworks. We are also increasingly focused on dual-use technologies that combine AI with climate solutions, cybersecurity, and urban innovation, ensuring Lisbon remains a living lab responsible for innovation at scale.

Europe's opportunity in AI will not be defined solely by technological capability, but by its ability to govern and scale innovation responsibly. If we succeed, cities like Lisbon can demonstrate that competitiveness and ethics are not opposing forces but mutually reinforcing ones.

To read about Unicorn Factory, please scan this QR code or visit this page.

THE FOMO TAX: HOW AI TOOL-CHASING BECAME THE NEW PRODUCTIVITY TRAP

Author
Ruslan Drozdov, Founder of Calma

There is a particular kind of busyness that feels like strategy. The morning scroll through AI launches. The group chats buzzing with tools that didn't exist last week. The product trials that consume an afternoon and disappear by Friday. For business leaders, this has become a familiar ritual — one that carries the satisfying weight of diligence and competitive awareness.

Ruslan Drozdov, founder of Calma — a performance lead generation company in the US market — knows this intimately. He built an entire system around it: curated feeds, alert channels, group chats designed to ensure nothing slipped through. What he discovered was that the system wasn't producing edge.

It was producing anxiety. In this opinion piece, Ruslan introduces the concept of the FOMO Tax — the hidden cost extracted not just in wasted hours, but in leadership attention burned on tools that never survive their own hype cycles. He makes the case that the operators genuinely ahead are not the ones testing the most tools, but the ones who chose fewer, went deeper, and built measurable processes around what they kept.

Something strange is happening to the focus of business leaders right now — and it doesn't look like a problem. It looks like diligence. It looks like staying sharp.

Every morning starts with a scroll: what dropped overnight, what launched, what everyone in the group chats is testing this week. Bookmarks that never get opened. Product trials that last two hours. Newsletter summaries of tools that didn't exist last Thursday. Founders and executives call this staying current. Competitive awareness. A professional obligation.

I called it that too. For longer than I'd like to admit.

At Calma, we run performance lead generation in the US market — an unforgiving environment of thin margins and high competition, where the difference between a good month and a bad one often comes down to cost-per-lead efficiency measured in single-digit dollars. When AI arrived as a genuine lever for compressing costs and accelerating work that used to require a much larger team, the window was open. We walked through it.

Before Calma, I spent years as an investor, watching founders navigate dozens of markets and conditions. None of those reps prepared me for a technology moving fast enough that the old lessons stopped applying — and the mistakes had no useful reference point.

Ruslan Drozdov
Founder of Calma

Because the stakes felt real, the logic followed: we couldn't afford to miss anything. If a competitor was running a better AI stack, we would feel it in the numbers before understanding why. Every week brought new tools claiming to change the economics of lead generation — smarter targeting, faster copy, cheaper automation. Some would be real. Some would be noise. The problem was there was no reliable way to tell the difference without testing everything. So we built a system to make sure nothing slipped through.

The information environment that emerged was designed to catch everything. Telegram channels, Reddit threads, Twitter feeds, group chats dedicated to AI releases — curated, maintained, set to alert. If something genuinely changed the game, you needed to know about it before your competitors did. That felt like edge.

What it actually produced was a very sophisticated version of anxiety.

The products we tested most frantically were the ones that disappeared fastest. The tools that made the biggest noise at launch were often the ones no one serious used three months later. The "game-changing releases" we dropped everything to try were, in most cases, forgotten by the people who had loudly announced them.

Meanwhile, the tools that actually changed how the business operates — the ones embedded in daily process, producing measurable results, used by the team without prompting — almost none of them came from that frantic monitoring. They arrived differently: through someone who had already used a tool for six months, through noticing what had survived the hype, through testing something quietly adopted by people who cared more about results than announcements.

The recognition came in one of those internal conversations where you stop optimizing the answer and just say it. The pattern, stated plainly: we are not fast enough to monitor and test every new product on the market, and this FOMO is costing us something real. What had felt like competitive vigilance was not a disadvantage we were managing — it was a tax we were choosing to pay.

We started calling it the FOMO Tax.

The real cost isn't the hours lost to testing. It's the cognitive overhead: the ambient low-level pressure of tracking a field that moves faster than any organization can absorb. Decision fatigue from evaluating tools that didn't need to be evaluated yet. Leadership attention — and team attention below it — spent on noise rather than signal.

No meaningful competitive advantage comes from it. And the trap is elegant: in a fast-moving environment, monitoring feels like action. It mimics productivity. It produces a sensation of forward motion without the risk or effort of actual building.

There is a way out of it — not by ignoring AI, which would be a different mistake, but by trading monitoring for measuring.

The shift involves three moves.

1. Audit what is genuinely embedded in your operation — not bookmarked, not trialed once, but running regularly and depended on by the team.

2. For each working tool, connect cost to output in the same unit: what it costs per month and what it replaces, expressed in money or freed capacity.

3. Draw the line most organizations stop short of — from saved hours to a number on the P&L: revenue, CAC, conversion.

A tool that has been running for ninety days without a traceable line to a number that matters is ambient spending, not investment.

There is a second payoff that almost no one talks about. Once the map is visible, the FOMO shrinks. You stop feeling behind because you can see what you have. The urge to chase the next launch fades when you realize you are not fully utilizing the stack you already built. Each of these moves opens its own set of questions, and I plan to keep working through them in longer form. The direction here is clear enough to start.

The operators genuinely ahead right now are not the ones who tested the most tools. They are the ones who picked fewer things, went deeper on them, and built processes around what they chose. They are not reading about AI — they are running it, iterating on it, and measuring what it returns.

Stop optimizing your information environment. Stop mistaking awareness for advantage. Measure what you have before you add anything new.

The tools that matter will still be there when the hype passes. The attention you spend monitoring the hype will not come back.

To read more about Calma, please scan this QR code or visit calma.io.

AI ETHICS AND INTEGRITY
INTERNATIONAL ASSOCIATION

UNITED FOR A BETTER AI FUTURE

To read more about our association, scan this QR code or visit www.ai-ei.org

Have an idea, collaboration proposal, or material to share? scan this QR code or visit www.ai-ei.org/contact

If you want to become a member, scan this QR code or visit www.ai-ei.org/membership

To subscribe to our updates, scan this QR code or visit www.ai-ei.org/newsletter

Thank you for reading AI Horizon Journal!

We hope this edition sparked ideas worth pursuing and conversations worth having. If you have any inquiries, partnership proposals, or would like to contribute to a future edition, we'd love to hear from you at info@ai-ei.org