AI ETHICS AND INTEGRITY
INTERNATIONAL ASSOCIATION

MARCH 2026

AI HORIZON JOURNAL

#3 EDITION

ARTIFICIAL INTELLIGENCE: MANAGING RISK WHILE UNLOCKING STRATEGIC VALUE

JOSÉ CARLOS SOLA OF AIJU ON WHY GOVERNANCE, NOT SPEED, IS THE REAL COMPETITIVE EDGE.

VDS 2025: GLOBAL INNOVATION MEETS ETHICAL GOVERNANCE

IN INTERVIEW WITH GRAZINA SAVICKAITE (ZADARMA) AND CLAUDIO GONZÁLEZ (BANSHEE TECHNOLOGIES)

HOW AI TOOLS MAY SUPPORT ARTISTS AND DESIGNERS IN CREATING HIGH-QUALITY CONTENT MORE EFFICIENTLY

IN CONVERSATION WITH DMYTRO MASLOV, GLOBAL AI LEAD AT PLARIUM

0 23991 18471 9

FEATURING VDS | ZADARMA | AIJU | PLARIUM | BANSHEE TECHNOLOGIES | PRANAVA INSTITUTE

CONTENTS

MARCH 2026 / ISSUE NO.03

INTRODUCTION

ASSOCIATION UPDATES

MEMBER'S VOICES

Member's Voices is a regular column where active members are invited to share insights, experiences, or initiatives related to ethical AI. If you are interested in contributing please reach out to us at info@ai-ei.org for more details.

COVER STORY

INDUSTRY INSIGHT

Artemiy Volodin
AIEI Vice President

WELCOME MESSAGE

We live in a strange moment. AI is everywhere in the headlines, yet still absent from most boardrooms in any meaningful way. The gap between the conversation and the reality has never been wider — and closing it is exactly what this journal exists for.

Edition #3 arrives at a time when the industry is, as one of our contributors puts it, somewhere between anger and acceptance. The hype cycle is maturing. The easy promises are wearing thin. The narratives that once filled conference stages — of frictionless automation, universal efficiency gains, and intelligence that practically deploys itself — are giving way to something more grounded, more honest, and ultimately more useful. What's left is the harder, more important work: governance that actually functions, tools that genuinely serve people, and frameworks built not for press releases, but for day-to-day use across organizations of every size and sector.

That shift is not a defeat. It is, if anything, the real beginning. The most consequential phase of any transformative technology is not when everyone is talking about it — it's when serious people start doing the quiet, unglamorous work of making it trustworthy. That is the phase we are entering now, and it is the phase this journal is built for.

Inside, you'll find voices who have moved well past theory. We look at Plarium, where AI earns its place by eliminating grunt work without touching the creative ceiling. We dive into Banshee Technologies, where data sovereignty is no longer a feature—it's a legal inevitability. From Zadarma's human-centric communication tools to AIJU's validation models, the theme is the same: no algorithm scales until it is understood.

And threading through all of it: the work of the AIEI Legal Committee, which has just launched its 2026 agenda with one clear mandate — turn principles into practice. Not more declarations. Not another set of guidelines that sit in a drawer. Working tools: a governance starter kit, a global regulatory portal, refined principles with real explanatory depth. The kind of infrastructure that lets organizations actually act on what they already believe.

This is not a publication about what AI might do someday. It's about what responsible actors are doing right now and what the rest of us can learn from them. Each edition, we try to bring you fewer buzzwords and more substance. Less speculation, more signal.

Read critically. Think practically. And if something in these pages challenges your assumptions, good. That's the point.

AI ETHICS AND INTEGRITY INTERNATIONAL ASSOCIATION

AI HORIZON CONFERENCE

WHERE RESPONSIBILITY MEETS THE FUTURE OF AI

At AI Horizon 2025, the conversation was less about what artificial intelligence can do — and more about what it should. Experts, business leaders, and government representatives gathered to explore ethical AI development, cross-industry collaboration, and the practical work of building systems people can trust.

The room was full of people who build things for a living — lawyers automating trademark filings, engineers rethinking the software development lifecycle, government officials designing regulatory sandboxes. But what united the day's speakers wasn't a shared enthusiasm for what AI can do. It was a shared unease about how it gets done.

AI Horizon 2025 opened with Sergiy Barbashyn, President of the AI Ethics and Integrity International Association (AIEI), setting a tone that would persist through every panel and pitch: the technology is moving faster than the institutions meant to govern it, and the gap is everyone's problem.

"Ethical AI practices have to keep pace with rapidly evolving laws — not trail them by five years."

Sergiy Barbashyn, AIEI President

Under Barbashyn's leadership, AIEI has spent the past year assembling infrastructure. A set of International AI Principles, drawn from twelve global sources and reviewed across eight countries, now offers companies a practical certification framework. A Legal Committee of over 22 experts tracks regulatory developments worldwide. A new journal circulates case studies and research across industries. None of it is glamorous. All of it is necessary.

AIEI — YEAR IN BRIEF

- International AI Principles: Consolidated best practices from 12 global sources, reviewed in 8 countries
- Legal Committee: 22+ experts tracking regulation and supporting AI deployment in law
- AI Horizon Journal: Platform connecting innovators and thought leaders across industries
- International Network: Cross-sector partnerships between startups, academia, and enterprise

AI IN THE ROOM WHERE IT HAPPENS

Ruba Hamam, Partner & Legal Director, at AÏP Genius, made a case that tends to get lost in broader AI debates: the real value often comes not from replacing experts, but from removing the friction that slows them down. Their platform handles trademark filings in the UAE in roughly ten weeks — against an industry standard of six to eight months. Renewals that once took days now take hours.

"We started with a narrow domain," she said, "built from verified data, and kept human judgment central."

Roba Hamam, Partner & Legal Director, AÏP Genius

That framing — narrow, verified, human-in-the-loop — recurred throughout the day. Igor Paniuk of Trinetix, drawing on over a decade working with Fortune 500 companies, described a spectrum of AI deployment strategies: from augmenting traditional software development cycles to fully AI-native systems where humans focus almost entirely on creative and strategic decisions. His advice was cautious in all the right ways: pilot on low-risk projects, assign internal champions for medium-impact work, apply context engineering only where reliability is non-negotiable.

"Piloting AI on low-risk projects first isn't timidity — it's the only strategy that scales."

Igor Paniuk, SVP of AI Strategy & Innovation, Trinetix

Ukraine's Ministry of Digital Transformation offered a different angle — what happens when the state itself becomes a testing ground. Their QA Sandbox program gives AI and blockchain startups free access to legal, business, and technical audits. Results stay confidential. The goal isn't oversight in a restrictive sense; it's the kind of structured support that turns a promising model into a compliant, deployable product. Backed by the UK-funded UK DIGIT initiative, the program positions Ukraine as a serious player in responsible AI development.

THE PANEL THAT DIDN'T STAY ON SCRIPT

The afternoon's most animated session was billed as "Responsible Innovation: Building Trust in the Age of Intelligent Systems." What the audience actually got was something more alive — a genuine debate about the present and future of AI practice, driven as much by the room's questions as by the panellists' prepared remarks.

The discussion ranged widely. What skills and mindsets will professionals need over the next five years as AI reshapes their industries? What does it actually mean to build trust in an AI system — and how do you measure it? Perhaps most unexpectedly, the panel explored the potential of cross-industry collaboration between GameDev and Cybersecurity to advance AI safety, arguing that disciplines with deep experience in adversarial thinking and complex system design have insights that more traditional AI ethics conversations often miss.

The lively back-and-forth with the audience made clear that these weren't abstract questions. People in the room were living them. Nobody resolved anything definitively — but that, from the energy in the space, seemed precisely the point.

PANEL PARTICIPANTS

- Catarina Farinha, AI Product Manager, Sword Health
- Professor Miguel de Castro Neto, Dean, NOVA Information Management School
- Dmytro Maslov, Global AI Lead, Plarium
- João Andrade, Partner, Hacken

STARTUPS, PRIZES, AND WHAT COMES NEXT

The Startup Pitch Session gave the floor to four ventures representing the next generation of AI innovation. Each brought a distinct take on applied AI — from intellectual property tooling to creative platforms — and each demonstrated the kind of domain-specific focus that the day's expert presentations had repeatedly identified as the path most likely to succeed.

Winners received prizes from SkyUp Airlines, while all conference attendees were offered special flight promotions. Partners, including Network VC were present throughout, underscoring their commitment to supporting innovation and building bridges across the AI ecosystem. The session was energetic, competitive, and a welcome reminder that the ideas driving AI's next chapter are still being formed by founders willing to bet on them.

STARTUP PITCH SESSION

- Marqea - Presented by Sandra Kobel
- Paintit.ai - Presented by Yulii Cherevko
- HostyAI - Presented by Joao Botelho
- SYLA - Presented by Mykola Lozinskyi

AWARDS AND CLOSING REMARKS

The conference concluded with an awards ceremony recognising 15 companies for significant contributions to ethical AI development. Among those honoured — AIP Genius, Hacken, Trinetix, and APDC — received their awards in person, a fittingly tangible acknowledgement of work that often happens far from the spotlight.

In his closing remarks, AIEI Vice President Artemiy Volodin struck a forward-looking tone, sharing three initiatives on the horizon: the launch of an international AI registry, the development of an AI assessment platform, and the opening of a new AIEI office in New York. Each represents a concrete step toward the kind of global infrastructure that ethical AI development needs — not just principles on paper, but systems and institutions capable of making them real.

AI Horizon 2025 reaffirmed its role as a platform for knowledge exchange, ethical practice, and cross-sector collaboration. The throughline of the day was clear: the future of AI is shaped not only by innovation, but by responsibility, transparency, and the willingness of an entire ecosystem to hold itself to a higher standard.

AIEI IN 2025: BUILDING THE INFRASTRUCTURE FOR RESPONSIBLE AI

From ethical principles and legal committees to a growing global network, AIEI spent 2025 turning ethical AI from a talking point into a working system.

A year ago, the conversation around AI ethics was still largely theoretical — a discussion of principles without the machinery to put them into practice. In 2025, the AI Ethics and Integrity International Association set out to change that. Bringing together professionals from law, technology, government, and international organisations, AIEI spent the year translating values into tools: frameworks that can be certified, committees that can advise, publications that can inform, and events that can connect. The result is an association that looks considerably different at the end of the year than it did at the beginning — one with more reach, more structure, and more to offer the professionals navigating AI's increasingly complex landscape.

A FRAMEWORK THE WORLD CAN USE

The centrepiece of AIEI's 2025 output is its International AI Principles — a document long enough to be comprehensive and short enough to be used. Drawing on twelve international sources and reviewed by ten experts across eight countries, the Principles were designed from the outset for global applicability, not just local relevance. They are publicly available on the association's website, and members can formally declare adherence, follow best practices, and receive certification demonstrating their commitment to responsible AI.

Certification matters here more than it might seem. In a field where "ethical AI" risks becoming a marketing phrase rather than a meaningful standard, a verifiable public commitment gives the declaration weight. Companies that go through the process aren't just endorsing a set of ideas — they are accountable to them.

LEGAL EXPERTISE ACROSS BORDERS

Regulatory fragmentation is one of the defining challenges of AI governance today. The European Union has produced some of the world's most demanding AI rules; the United Kingdom has taken a more permissive approach; the United States still lacks a coherent federal framework. For any organisation operating internationally, navigating this patchwork is a significant burden — and the cost of getting it wrong is rising.

AIEI's Legal Committee was established to address exactly this. Drawing representatives from seventeen countries, the committee supports members in understanding and applying legal and regulatory requirements across jurisdictions. Its initial priorities reflect the breadth of what's needed: reviewing and updating AIEI's Declarative Principles 2.0 to keep them clear and sector-agnostic; developing an AI Governance Starter Kit — a set of adaptable policy templates that companies can deploy without building from scratch; and tracking AI-related legislation globally so members aren't caught off-guard by developments in markets they operate in.

Beyond its internal work, the committee contributes to webinars, seminars, and professional events — sharing knowledge outward and building AIEI's visibility as a credible voice in the international policy conversation.

KNOWLEDGE THAT TRAVELS

The AI Horizon Journal emerged in 2025 as AIEI's platform for translating expertise into accessible insight. It publishes interviews, case studies, and articles from practitioners across the AI field — contributors have included the Paris Bar Association, Ukraine's Ministry of Digital Transformation, and leading companies working at the frontier of applied AI.

For members, the journal offers something specific: real-world examples of AI implementation drawn from different sectors, alongside practical guidance on ethical deployment and governance. It also gives visibility to approaches that rarely surface in mainstream coverage — from startups experimenting at the edges of the field, to research institutions rethinking foundational questions, to international organisations working to harmonise standards across borders. The cumulative effect is a publication that treats AI governance not as a compliance problem, but as a discipline worth taking seriously.

A CONFERENCE THAT MEANS

AI Horizon 2025 was, by most accounts, the year's most concentrated gathering of people who actually build AI systems for a living. Experts, business leaders, startup founders, and government representatives spent the day moving between keynotes, panels, and pitch sessions — not to celebrate the technology's potential, but to interrogate how it should be developed and deployed.

The conversation was practical throughout. Presentations covered AI's transformation of intellectual property management, the changing nature of software development, and the role governments can play in supporting responsible innovation without stifling it. A panel on trust and responsible innovation drew out questions the field is still working through: what skills will professionals need as AI reshapes their industries? What does cross-sector collaboration between, say, cybersecurity and game development actually look like in the context of AI safety? The audience's energy suggested these weren't abstract concerns.

The day closed with an awards ceremony recognising fifteen companies for meaningful contributions to ethical AI — a reminder that the field already has examples worth pointing to.

PRESENT IN THE WORLD

An association that only speaks to its own members has limited reach. In 2025, AIEI representatives attended more than nine conferences across seven countries — sharing best practices, presenting the association's initiatives, and expanding its professional network beyond the communities it already serves. Each appearance was an opportunity not only to raise awareness of AIEI's work, but to bring back perspectives from practitioners operating in very different regulatory and cultural contexts.

LOOKING TO 2026

The next year's agenda reflects an association that has moved from building foundations to scaling what sits on top of them.

A third AIEI base, to be opened in New York, will establish a meaningful foothold in the American market — one of the world's most active AI ecosystems, and one where AIEI's voice has room to grow. A free and open International AI Register will allow companies and organisations to make their AI activities and commitments visible on a global stage, creating a searchable record of who is doing what and to what standard.

Two new platforms round out the plan. An AI Ranking Application will evaluate and compare achievements across AI sectors — recognising successful projects and encouraging the kind of peer comparison that drives improvement. And an AI Assessment Platform will give companies and experts structured guidance on legal and regulatory requirements across jurisdictions: the EU's strict rules, the UK's softer framework, and the still-unresolved landscape in the United States. For organisations trying to operate responsibly across borders, this kind of clarity is not a luxury.

Taken together, the plans for 2026 suggest an organisation that has earned the right to think at scale — and that understands the difference between talking about responsible AI and building the systems that make it possible.

AI ETHICS AND INTEGRITY INTERNATIONAL ASSOCIATION

AIEI LEGAL COMMITTEE
LAUNCH OF WORK AND 2026 AGENDA

The Legal Committee of the AI Ethics and Integrity International Association (AIEI) has officially commenced its work for 2026. Following its inaugural meeting, the Committee has established a strategic foundation aimed at one of the industry's most pressing challenges: moving beyond high-level debate toward practical, scalable governance.

The session brought together a diverse group of practitioners and experts across law, artificial intelligence, and AI governance. From the very beginning, it was clear that the synergy within the Committee is "not merely theoretical, but practical and scalable," reflecting a collective desire to produce results that resonate across the professional community.

A VISION OF APPLIED INTEGRITY

The meeting opened with remarks from AIEI leadership, who emphasized the critical role the Legal Committee plays in strengthening trust in the Association, its principles, and its practical instruments. The mission is not simply to discuss the abstract nature of regulation, but to develop "working solutions that can be applied by businesses, legal teams and professionals across different levels of expertise."

Taras Lytovchenko, Head of the Legal Committee, highlighted the necessity of this hands-on approach.

"We see how fast AI law is moving and how hard it is for teams to keep up across jurisdictions," Lytovchenko noted. "What stood out in our first meeting was the mix of perspectives and the shared focus on outcomes."

REFINING THE DECLARATIVE PRINCIPLES 2.0

A primary task identified by the Committee is the review and update of the AIEI's core principles. As the structural framework of the Association, these principles must remain agile. As the Committee noted during the session, "if they are too general, they become impractical; if outdated, they risk losing credibility."

To ensure these principles remain a "living" resource, the Committee will:

1. Review approved principles to identify gaps or unclear provisions.
2. Refine wording to enhance practical applicability for day-to-day operations.
3. Add concise explanations to ensure accessibility for professionals in other disciplines, such as engineering and procurement.

Additionally, the development of white papers is planned to demonstrate how these principles can be applied at different levels of business, including internal processes and vendor management.

AI GOVERNANCE STARTER KIT

A second key focus area is the development of the AI Governance Starter Kit. This initiative is designed as a structured set of policies, procedures, and practical guidance that companies can adopt and adapt to their specific needs.

The Committee clarified that the Starter Kit is "not a one-size-fits-all solution, but a flexible framework that can be tailored to various business models, organizational sizes and regulatory environments." Its aim is to provide concrete tools for implementation, helping organizations translate high-level ethics into actionable measures and empowering them to "implement ethical, responsible and legally compliant AI practices globally."

REAL-TIME MONITORING: THE GLOBAL PORTAL

To address the fragmented nature of international AI law, the AIEI plans to establish a dedicated Global Portal. This will be a centralized resource where lawyers and experts from different jurisdictions can provide real-time updates on regulatory developments in their respective countries.

This portal will enable members to follow global trends and changes in legislation as they happen, supporting informed decision-making and fostering a shared understanding of governance standards across different industries and jurisdictions.

NEXT STEPS AND FORWARD MOMENTUM

The meeting concluded on a positive and forward-looking note. Following the session, the Committee agreed to identify volunteers for priority workstreams and form specialized subgroups to handle the technical aspects of the Starter Kit and the Global Portal.

Beyond internal coordination, the Committee will participate in seminars, webinars, and international conferences to ensure the AIEI's voice remains central to the global conversation.

LEGAL COMMITTEE COMPOSITION

The current composition of the Committee reflects a unique combination of practitioners and experts capable of delivering meaningful results for both the professional community and the Association as a whole.

> *"We are not here only to comment on regulation. We are here to help organizations implement governance that works in real life, across different sizes and legal systems. That is why our first priority is to make the AIEI Principles easier to use in day to day work, and why we are building the AI Governance Starter Kit and baseline controls to turn principles into action and help the community stay current as the rules evolve."*
>
> – Taras Lytovchenko, Head of Legal Committee

COMMITTEE LEADERSHIP

Taras Lytovchenko
Chair

Dr. Edvinas Meskys
Deputy Chair

Roba Hamam
Deputy Chair

FULL MEMBERS

Roman Zavrsek

Oleksii Shamov

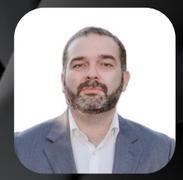

Vitor Neves

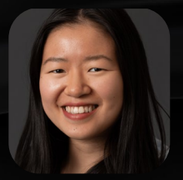

Janis Wong

Nélitié Zingoua Kouadio

Andriy Barbashyn

Hanna Bondarenko

Alexandr Chernykh

Olha Osinska

Amisha Mittal

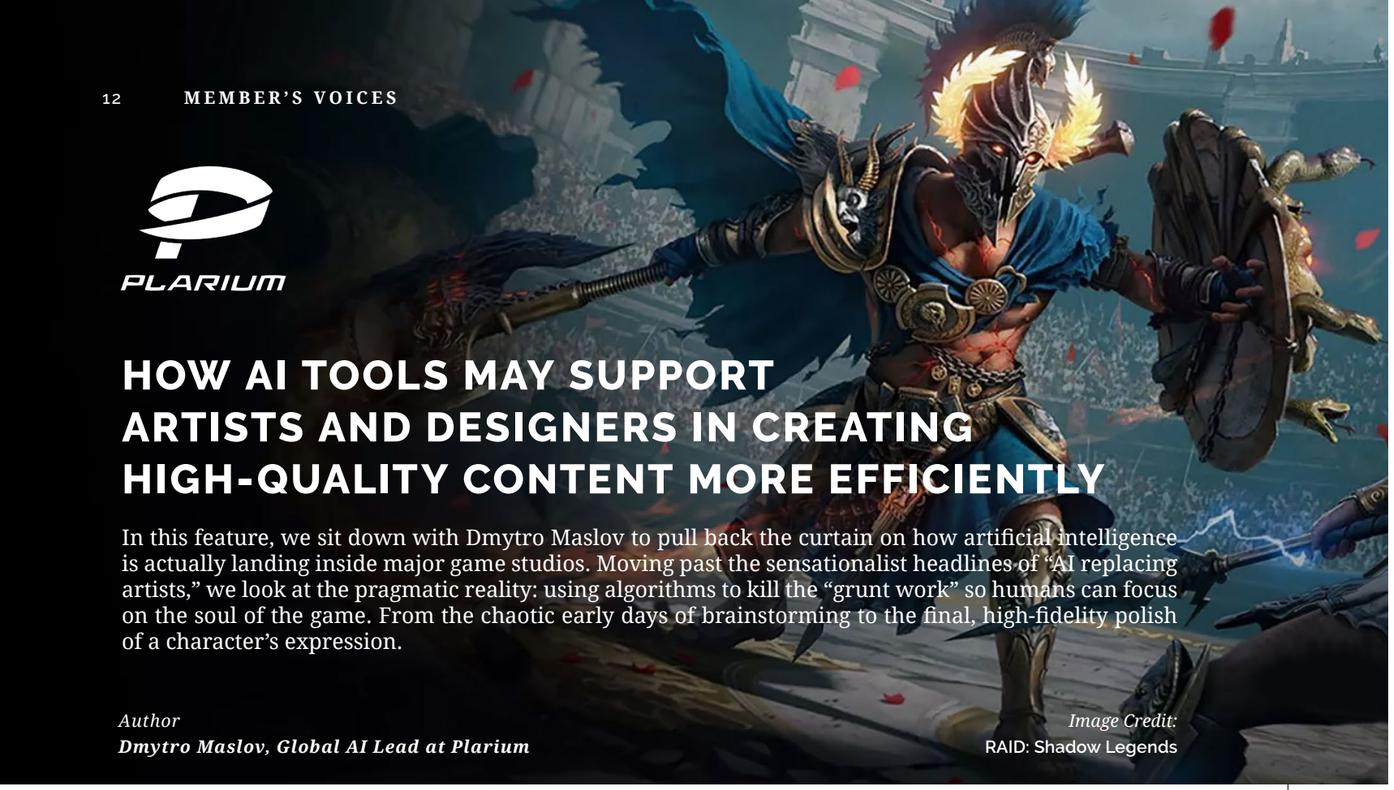

HOW AI TOOLS MAY SUPPORT ARTISTS AND DESIGNERS IN CREATING HIGH-QUALITY CONTENT MORE EFFICIENTLY

In this feature, we sit down with Dmytro Maslov to pull back the curtain on how artificial intelligence is actually landing inside major game studios. Moving past the sensationalist headlines of "AI replacing artists," we look at the pragmatic reality: using algorithms to kill the "grunt work" so humans can focus on the soul of the game. From the chaotic early days of brainstorming to the final, high-fidelity polish of a character's expression.

Author
Dmytro Maslov, Global AI Lead at Plarium

Image Credit:
RAID: Shadow Legends

ABOUT PLARIUM

Plarium is an international video game developer and publisher, founded in 2009, with studios across Ukraine, Israel, Finland, Poland, and Spain. Its portfolio includes titles such as Raid: Shadow Legends, Mech Arena, and Merge Gardens, available on iOS, Android, Steam, Epic, and Plarium Play — the company's own PC/Mac publishing platform, which features both Plarium titles and third-party titles such as Heroes of History from InnoGames.

ABOUT AUTHOR

With a career in technology spanning over 15 years, Dmytro began as a backend developer before progressing through a range of roles including frontend engineer, team lead, solution architect, and technology manager — balancing development, team leadership, and the management of technical solutions throughout.

For the past decade, he has been at Plarium, where he currently serves as Global AI Lead, responsible for implementing artificial intelligence across all of the company's studios worldwide.

Some people think AI does more than it actually does when it comes to making visual content. AI's main advantage here is increased efficiency, not the replacement of human skill.

Like many fields, content creation involves many routine tasks. Right now, AI mainly helps us speed up these parts of the process.

There are many situations where AI can be used in asset creation: they can be found directly in the game, for example, icons, avatars, splash screens, event banners, etc. Along with different marketing materials outside of the game itself. Creating content with AI isn't as simple as writing a prompt and using the first image it gives you. Even high-quality images need editing, and sometimes explaining every detail or nuance to AI takes longer than doing it manually. AI doesn't replace an artist's skill or intuition, especially in tasks that need a creative or human touch.

To sum up, AI is most useful in three phases of our content creation process.

Phase 1: Brainstorming and Testing Ideas. The first phase is brainstorming and testing ideas, such as creating mood boards or storyboards. A lot of this work gets thrown out, so it's okay if AI makes mistakes, like drawing six fingers or getting the perspective wrong.

Phase 2: Supporting Artists' Work. For 3D Pipelines, for instance, AI can help by generating the ba- sic shape, so artists don't have to sculpt it from scratch. After that, artists can use their usual tools, like Blender, to finish the job.

Phase 3: Polishing. The last phase is polishing. Here, you already have your own art, and AI can help add details, like changing the lighting, adding textures or scuffs, or adjusting the emotion in a character.

WHY HUMAN CREATIVITY AND INTUITION REMAIN AT THE CORE OF PLARIUM'S DEVELOPMENT PROCESS

Many factors make games engaging, such as their mechanics and visual style. Great visuals and engaging mechanics keep players coming back, while fresh updates give games new life—all made possible by the creativity and skill of my colleagues.

Plarium's artists are very talented. Just look at the unique style of RAID Shadow Legends! Game designers create features and mechanics, and developers bring them to life in the game. Managers and coordinators handle lots of information and help different teams work together. There isn't enough space here to describe every role and its impact, but everyone makes a huge difference in creating our games.

This teamwork is the foundation of game development, not just at Plarium but across the industry. People make games. Their creativity and intuition are essential for success.

That's why creativity and intuition have always been at the heart of our work. They help us succeed, and players have enjoyed our games for over 10 years. That speaks for itself.

TRENDS AND CHALLENGES AI BRINGS TO THE GAME INDUSTRY

Everything below is my personal opinion. I do not claim to be 100% objective.

When it comes to trends and challenges, the gaming industry is a lot like other software development fields. Games are products built with code, and they also require graphics, marketing, post-release support, bug fixes, and more.

Of course, games have some unique challenges. Their graphics are more complex, there's sound design, and developers have to focus on things like optimization and bundle size.

With that in mind, here are some trends I see in game dev, beyond the ones that apply to other software fields:

AI is finding its place in the game dev industry, especially in speeding up pre-production tasks like brainstorming and testing ideas. Product managers and game designers can now try out concepts on their own, without needing developers right away. I think the current level of 'vibe coding' is good enough for this.

Another interesting trend is "world models", which are getting a lot of attention online. For example, Google's DeepMind recently showed a new version of "Project Genie", which can create interactive virtual worlds in real time. Now, you don't have to build assets or levels in Unity just to test ideas. With your creativity and some prompts, AI can handle the rest, saving the team significant time and effort.

Dmytro Maslov
Global AI Lead at Plarium

As for the challenges. You might know the five stages of accepting something inevitable: denial, anger, bargaining, depression, and acceptance. I'd simplify it to just three: denial, anger, and acceptance.

Right now, the gaming industry is somewhere between anger and acceptance. Most people agree, and I do too, that AI is here to stay. All signs point to that.

We've moved past the stage where everyone joked about how awkward AI was. Still, many players aren't happy about AI in games. Sometimes, people confuse using AI as a tool for creators with simply dropping AI-generated assets into games without any changes.

I think the best solution is to give it time. AI is becoming a normal part of our lives, and soon most people won't pay much attention to it. While we wait for this shift, it's important to highlight the potential benefits of AI, such as faster game releases, greater variety, and lower development costs. These benefits aren't fully here yet, but they show what's possible. Clear communication is key. Not everyone wants to hear about it now, since hate is trendy nowadays, but most players don't post on social media and outnumber the critics. We should focus on reaching these players.

A VISION FOR THE NEXT FEW YEARS OF AI-POWERED ENTERTAINMENT

In 2026, AI will shift from impressive demos to creating real value. That's what leaders in Big Tech and AI researchers believe. We're past the early development stage, and I hope we're ready to find practical ways to use AI effectively in our work and daily lives.

AI probably won't change overnight from a helpful companion tool to something completely new, but it will become a regular part of our devices and routines, just like other technology has.

I think AI will become even more common and integrated in the next few years. After that, we'll be ready for the next wave of technology, with its own trends, challenges, and questions.

VDS 2025

WHERE GLOBAL INNOVATION MEETS ETHICAL GOVERNANCE

In 2025, representatives from the AI Ethics and Integrity International Association (AIEI) attended the VDS 2025 conference in Valencia. The AIEI team connected with a wide range of startups, corporations and thought leaders, engaging in dynamic discussions on how emerging technologies, particularly artificial intelligence are shaping the future of innovation across industries.

Conversations explored key topics such as the ethical implementation of AI, the role of AI in smart cities and data governance and the challenges of integrating rapid technological advancements with societal and regulatory frameworks. The delegation gained valuable insights from influential speakers and innovators, including leaders from HomeDock OS, Zadarma and AIJU, who are driving the next wave of technological solutions.

VDS 2025 is one of Europe's premier technology and innovation forums, attracting thousands of participants from across the globe. Founded with the mission to foster collaboration, showcase cutting-edge solutions and promote sustainable technological growth, VDS provides a platform where startups, investors and established companies converge to exchange knowledge, spark partnerships and shape the trajectory of future tech ecosystems. Through this environment, AIEI deepened its understanding of emerging trends and explored opportunities to advance responsible AI practices in alignment with global innovation priorities.

RESPONSIBLE BY DESIGN: ZADARMA'S APPROACH TO AI IN COMMUNICATION

Zadarma is a global Voice over Internet Protocol (VoIP) provider with 19 years of experience in electronic communications. With its in-house development team, Zadarma continuously evolves by creating new services integrated with telephony, including Cloud PBX, CRM, and AI voice agents. The company focuses on providing reliable and affordable telephony and AI-based communication solutions.

What are the key ethical principles underlying your work with AI?

At Zadarma we base our AI ethical principles on recommendations from organizations such as UNESCO, the EU AI Act, and others. The key principles are:

Privacy and data protection: We prioritize user privacy and comply with regulations such as GDPR.

Transparency and explainability: We provide a clear guide to setting up the AI agent. All answers and decisions made by the agent are preconfigured by the user.

Human agency and oversight: We ensure that AI systems improve, rather than replace, human participation and decision-making.

Safety, reliability, and security: Our AI voice agent is reliable, safe, and secure, with multiple security layers including two-factor authentication and automatic breach detection. Our infrastructure and workflows have been independently verified to meet the international information security standard ISO/IEC 27001.

Inclusiveness: We make AI technologies accessible to all users, independent of their background.

How do you ensure transparency and keep users informed about the functioning of your AI agent?

The AI agent comes with a step-by-step configuration guide that clearly explains each setup step and its meaning. Users can build the agent themselves by selecting from the provided options and adding their own knowledge base. Our 24/7 multilingual customer and technical support is also available to assist with any questions.

Grazina Savickaite
Operations and Content Specialist

How does your AI voice agent change the way users interact with the company's services?

Our AI agent helps companies reduce employee workload, allowing teams to focus on complex tasks and resolve customer issues faster and more efficiently. The agent can also process calls outside working hours, which helps businesses grow their customer base. By automating routine tasks such as answering basic questions, it improves both customer communication and employee satisfaction.

What new AI development directions does your company plan to pursue in the coming years?

Our future development will focus on real-time AI to reduce response latency and improve conversational quality. Our goal is to improve agent understanding and performance, and expand support for multiple large language models, including Gemini.

That said, there is no magic button that configures the AI agent automatically. For the agent to work reliably, users should provide accurate information and clear rules — this avoids the agent guessing answers when information is missing from the knowledge base.

How does your company combine innovation with ethical responsibility in its daily AI work?

We combine innovation with ethical responsibility by integrating ethical principles directly into the development of our AI systems. We protect user privacy, regularly audit our models for reliable performance, and maintain a pace of fast innovation — while staying transparent, accountable, and trustworthy for our users.

YOUR DATA, YOUR RULES: INSIDE HOMEDOCK OS

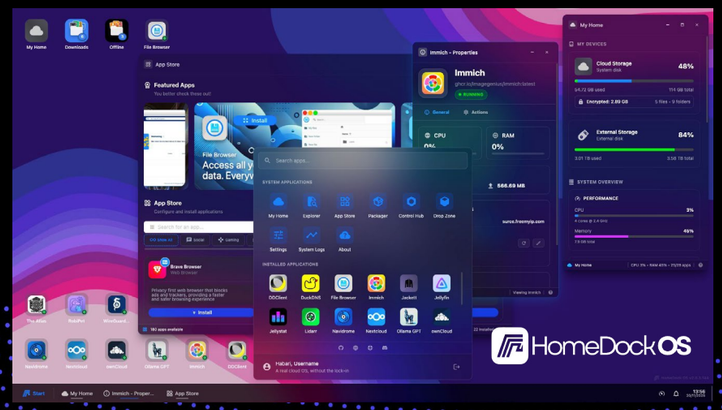

Banshee Technologies develops HomeDock OS, a cloud operating system that turns any computer into a private cloud server with over 200 self-hosted applications. HomeDock OS Enterprise brings local AI, end-to-end encryption, and EU-compliant data sharing to organizations that refuse to compromise on data sovereignty.

As data privacy regulations tighten and enterprises grow wary of cloud dependency, a new generation of self-hosted infrastructure is emerging. In this interview, Claudio González, CEO of Banshee Technologies, explains how HomeDock OS is positioning local AI and private cloud as a practical, compliance-ready alternative to mainstream cloud services — and why he believes this shift is no longer a preference, but a legal inevitability.

How do you plan to achieve data privacy?

We built HomeDock OS around one principle: your data should stay yours. The platform follows an open-core model — anyone can use it on our managed cloud or install it offline on Windows, macOS, or Linux. Our Enterprise modules — Local GPT, Secure Tunnel, Data Space — are exclusively on-premise. That's intentional. If a bank needs AI that never phones home, or a government agency needs encryption beyond HTTPS, those tools should only run on hardware they physically control. No telemetry, no cloud dependency.

What is the concept of Local AI for digital services?

Local AI means running language models directly on your own servers using open-source, quantized models. No cloud subscriptions, no data leaving your network — ever. Our Local GPT module runs models ranging from lightweight 0.5B-parameter on modest hardware to advanced 20B-parameter reasoning models. For most business tasks you do not need a trillion-parameter cloud model. A well-chosen local model delivers practical results with zero privacy risk. We've also trained our own model, HomeDock OS 3B Instruct, fine-tuned on our entire documentation — an open-source built-in copilot, available on Hugging Face.

Which digital services can be automated using Local GPT?

With HomeDock OS 2.0 we introduced Prism Window Manager — a full desktop environment in the browser with multitasking, window snapping, and everything you'd expect from a real OS. Local GPT lives inside it as a native application. Organizations use it for code generation, document drafting, technical analysis, data summarization, and internal Q&A, with full export options. If a conversation is too sensitive even for local storage, Privacy Mode keeps everything in memory only — nothing touches the disk.

Will "Self-hosted AI" become the standard for dev teams in 2026?

The shift is already underway. When AI runs at the OS level alongside Docker containers and CI pipelines, developers can query models against their actual codebase without uploading proprietary code to external services. For regulated industries this is not optional — it is a compliance requirement. We expect 2026 to be the year dev teams start treating local AI as the default for anything involving sensitive code. The momentum around projects like OpenClaw is a clear signal: not just for basic queries, but for full AI agents with complete access to the machine. We plan to bring it to our App Store once their security model matures.

What share will private local AI servers have compared to public clouds in five years?

Public clouds aren't going anywhere for consumer services and large-scale training. But for enterprise inference, we see 30 to 40% moving to private infrastructure by 2030. Open-source models are getting remarkably good and remarkably fast. Local deployment is becoming less of a preference and more of a legal reality. That's where HomeDock OS lives. We want running your own private cloud and local AI to feel as simple as signing up for a SaaS product. We think we're getting there — closing the gap.

Claudio González
CEO at Banshee Technologies

ARTIFICIAL INTELLIGENCE: MANAGING RISK WHILE UNLOCKING STRATEGIC VALUE

Author
José Carlos Sola,
Head of the Integrated Digital Technologies (IDT) Area at AIJU

As artificial intelligence becomes increasingly embedded in business strategy, the gap between ambition and operational reality continues to widen. In this article, José Carlos Sola, Head of the Integrated Digital Technologies area at AIJU, sets out a practical framework for AI adoption — one built on governance, measurable results, and human oversight rather than hype. Drawing on real-world initiatives across industry, healthcare, and education, he makes the case that responsible AI is not a constraint on innovation, but its most reliable foundation.

Artificial intelligence has rapidly become embedded in industrial discourse, often accompanied by inflated expectations and narratives that are more aspirational than operational — narratives that do not always translate into tangible organizational impact. Mindful of this reality, at AIJU we promote an evidence-based, results-oriented approach: a model that combines technological ambition with methodological rigor, placing responsibility and governance at the core of strategy.

Within this framework, artificial intelligence, including generative AI, is not conceived as a universal solution or as a technology for automatic deployment. Its true potential emerges when it addresses a clearly defined need, integrates coherently into existing processes, and operates under qualified human oversight. Absent these conditions, risks increase significantly: low value-added solutions, unnecessary technical complexity, or misalignment with organizational strategy.

"Competitive differentiation does not lie in adopting AI first, but in integrating it through a clear governance model, impact metrics, and strategic alignment."

This vision is articulated in our guide, "IA Hacia una industria inteligente" ("AI Toward an Intelligent Industry"), a document that systematizes practical and operational criteria for secure, effective implementation fully aligned with the real needs of the business ecosystem.

FROM VALIDATION TO SCALE: WHEN TO EXPAND AI IMPLEMENTATION

At AIJU we apply a key principle: before scaling, we validate and thoroughly understand the technology. Adoption is structured through pilot projects, rigorous validation of results, and continuous refinement. Algorithms are not implemented merely because they are technologically available, but when they demonstrate clear and quantifiable improvements in efficiency, information quality, or time reduction.

This approach makes it possible to identify precisely where AI generates sustainable competitive advantages and where a more cautious stance is warranted. The automation of repetitive, structured, high-volume tasks typically delivers demonstrable returns. By contrast, delegating complex decision-making without human supervision continues to present technical, regulatory, and reputational risks that must be managed prudently.

"Scaling without measurement, learning, and adjustment is not innovation — it is exposure to risk."

Accordingly, we conceive artificial intelligence as a tool for amplifying expert knowledge rather than replacing professional judgment. This pragmatic approach has proven decisive in integrating AI sustainably across sectors, ensuring gradual adoption aligned with each organization's operational realities.

ENSURING ETHICAL, SECURE, AND REGULATORY-ALIGNED AI IMPLEMENTATION

Beyond its efficiency or innovation potential, accountability constitutes a structural pillar of any artificial intelligence strategy. Issues such as data protection, compliance with the General Data Protection Regulation (GDPR), and human oversight at critical control points are not treated as formalities. Instead, they are enabling conditions for technological viability in real-world environments, particularly in regulated or mission-critical sectors.

Image Credit:
Laboratory at at AIJU Technological Institute for Children's Products & Leisure.

aiju TECHNOLOGICAL INSTITUTE FOR CHILDREN'S PRODUCTS & LEISURE

"Trust is the true strategic asset in any AI deployment. Without robust ethical frameworks, data traceability, and effective human oversight, potential benefits quickly erode."

AIJU actively contributes to defining and implementing best practices for the responsible use of artificial intelligence, with particular attention to data governance, system transparency, and the organizational impact of algorithmic decision-making. This approach extends to every domain in which AI is becoming increasingly relevant, ensuring ethical, secure, and fully compliant use under applicable regulatory frameworks.

IMPLEMENTING THE STRATEGY IN PRACTICE

This strategic model has been realized through diverse initiatives in which artificial intelligence is not an end, but rather an enabler for addressing specific challenges.

In the industrial domain, projects such as AI4TOYS and AI4VET have applied AI techniques to enhance design, manufacturing, and maintenance processes, combining data analytics, computer vision, and machine learning. The objective has been to increase efficiency and reduce error rates by embedding AI within existing workflows under conditions of continuous monitoring and validation.

In the healthcare sector, GLIO-IA demonstrates how image-analysis algorithms can support medical

José Carlos Sola
Head of the IDT Area at AIJU

professionals in diagnosing brain tumors. AI does not replace clinical decision-making; rather, it facilitates analysis and provides additional information that enhances diagnostic accuracy, while ensuring that the specialist remains central to the process.

Education represents another field of practical application. Through EDU4AI, machine learning–based solutions have been developed to adapt educational content and foster digital competencies, avoiding generic approaches in favor of realistic and measurable personalization.

Advanced data management and traceability also illustrate this model in practice. DLT4AITOYS and AIPASSPORT-GUARDNET combine artificial intelligence and blockchain technologies to develop digital product passports — AI enables the analysis of large data volumes and the detection of relevant patterns, while blockchain ensures data integrity and reliability, facilitating compliance with regulatory and sustainability requirements.

The experience accumulated through these initiatives confirms a strategic conclusion: artificial intelligence is neither an automatic nor a universal solution. However, when applied judiciously, under robust governance frameworks and with a clear results-oriented focus, it can become an effective catalyst for process improvement and the sustainable transformation of the business ecosystem.

The Pranava Institute
Center for Emerging Technology and Policy

SHAPING RESPONSIBLE TECHNOLOGICAL FUTURES FROM INDIA

Author
Shyam Krishnakumar
Executive Director and Co-Founder,
The Pranava Institute

Author
Amisha Mittal
Research Analyst,
The Pranava Institute

REFERENCES & FOOTNOTES

1. Epistemic Inequalities and Decoloniality: Mignolo, Walter D., The Darker Side of Western Modernity: Global Futures, Decolonial Options (Duke University Press 2011).
2. Global AI Cultures: Natale, Simone, Federico Biggio, Payal Arora, et al., "Global AI Cultures" (2025) 68(9) Communications of the ACM 37. Available at: https://doi.org/10.1145/3722547.
3. WEIRD Contexts in AI: Septiandri, Ali Akbar, Marios Constantinides, Mohammad Tahaei, and Daniele Quercia, "WEIRD FAccTs: How Western, Educated, Industrialized, Rich, and Democratic is FAccT?" (ACM Conference on Fairness, Accountability, and Transparency, Chicago, 12-15 June 2023). Available at: https://doi.org/10.1145/3593013.3593985.
4. AI Safety and Ethics: Public Citizen, "Counterfeit Companionship: Big Tech's AI Experiments Sacrifice Safety for Profit." Available at: https://www.citizen.org/article/counterfeit-companionship-big-tech-ai-chatbots/.
5. Legal Challenges and Mental Health: Ummer-Hashim, Shabna, "AI Chatbot Lawsuits and Teen Mental Health" (American Bar Association, 27 October 2025). Available at: https://www.americanbar.org/groups/health_law/news/2025/ai-chatbot-lawsuits-teen-mental-health/.
6. AI and Psychological Harm: Los Angeles Times, "Lawsuit alleges Google chatbot was behind a user's delusions and death" (6 March 2026). Available at: https://www.latimes.com/business/story/2026-03-06/lawsuit-alleges-google-chatbot-was-behind-users-delusions-death.

The Pranava Institute (TPI) is a New Delhi based research organisation working at the intersection of emerging technology, society, policy, and design to shape sustainable technological futures. Our work focuses on understanding how technological systems interact with institutions, markets, and communities, and on developing practical frameworks that support responsible and inclusive technological futures.

BRIDGING RESEARCH ON TECHNOLOGY, SOCIETY, POLICY AND DESIGN

TPI's work spans two verticals: Digital Economy and Tech Geopolitics, as well as Technology, Society, and Design. Our areas of work include Responsible Deployment of Artificial Intelligence (AI) in the Public Sector, Governance of Digital Public Infrastructures (DPIs), Critical Minerals Supply Chains, Semiconductors and Electronics Manufacturing, Trust and Safety Online, and Youth and Digitalisation. We believe in building on India's unique social, cultural and epistemic context to shape emic technological futures.

INFORMING TECHNOLOGICAL POLICY THROUGH ETHICS FOCUSED RESEARCH

Ethical considerations are often discussed within technology debates, yet translating principles into practice remains a major challenge. Our research focuses on operationalising ethical frameworks so that policymakers, developers, and institutions can apply them in real-world contexts.

• Youth Digital Cultures Lab: We recently launched the Youth Digital Cultures Lab (YDCL), a project supported by the Include+ Network at the University of Leeds. YDCL integrates ethics into technology governance by grounding discussions about digital technologies in lived experiences, cultural diversity, and participatory research. The lab abstains from approaching ethics solely through top-down regulatory frameworks. It focuses on how technologies are experienced by young people in everyday life and how those experiences can inform more responsible policy and design practices. The main questions explored within the lab are centred around themes of meaningful digital inclusion, digital well being, and non-western futures.

• Unboxing Tech Toolkit Series: Our Unboxing Tech Toolkit (UTT) engages ethics in technology by helping young people examine the often invisible infrastructures and value systems embedded within digital technologies. The toolkit has three interactive modules to unpack different ethical questions. The first module – "The Design Did It, Not Me!" helps youth recognise how platform design can influence their attention and behaviour; the second module - "The Materiality of the Smartphone" traces the hidden global supply chains behind everyday devices; and the third module – "Unboxing Internet Infrastructures" explores the environmental costs of the internet. The toolkit turns complex issues like deceptive design, sustainability, and digital responsibility into hands-on learning activities. UTT situates everyday digital practices within broader ethical debates about environmental responsibility, responsible innovation, and participatory technology governance.

• Responsible AI in the Public Sector: Our Responsible AI in the Public Sector project seeks to translate existing ethical guidelines into practical fitness checks that government decision-makers can use before deploying AI systems in public services. By providing accessible tools and policy recommendations, the project aims to support responsible AI adoption while mitigating potential harms.

BUILDING NON WESTERN, DIVERSE DIGITAL FUTURES WITH YOUNG PEOPLE

AI today raises pressing ethical challenges that extend beyond technical performance and into questions about knowledge, power, and human well-being. One significant concern arises from the persistence of historical and epistemic inequalities embedded within digital systems. AI systems risk reproducing these epistemic inequalities because they are often trained on datasets and frameworks shaped by dominant cultural contexts, particularly Western institutions. Such dominant ways of thinking continue to influence how knowledge, education, and governance are structured. Furthermore, many digital systems assume users originate from "WEIRD" (Westernised, Educated, Industrialised, Rich, Developed) contexts, which can marginalize users from different linguistic, cultural, or socio-economic backgrounds. Ethical AI governance must therefore move beyond abstract principles and interrogate how design processes themselves structure power relations and shape the lived experiences of diverse communities. The Youth Digital Cultures Lab (YDCL) at The Pranava Institute seeks to engage with these questions by foregrounding youth perspectives and culturally situated experiences of technology.

In this sense, the ethical challenge is not merely bias in datasets but the deeper reproduction of global knowledge hierarchies, where AI may marginalize indigenous knowledge systems or non-Western epistemologies by encoding a narrow conception of legitimate knowledge.

BUILDING EVIDENCE THAT INFORMS SAFE AND ETHICAL AI

As AI tools become more conversational and responsive, they increasingly simulate empathy and companionship. Recent lawsuits involving Character.AI and Google's Gemini allege severe psychological harm linked to emotionally intense chatbot interactions, including cases involving minors and users in vulnerable mental states. Our research project Feeling Automated examines the societal and legal implications of such systems, particularly the indirect risks associated with emotional manipulation, psychological dependency, and blurred boundaries between human and machine interaction.

DOING OUR BIT TO BUILD A MORE ETHICAL WORLD FOR TOMORROW

Responsible innovation requires integrating ethical reflection at the earliest stages of technological development and engaging meaningfully with stakeholders through participatory approaches and co-creation processes. This means involving communities, youth, policymakers, and affected users in shaping technologies. Technologies should be designed to be adaptable to different users and contexts. Achieving this requires diverse design teams, diverse datasets, co-creation with users, and flexible interfaces. Interdisciplinary collaboration across law, social sciences, design, public policy, and other domains can further help anticipate unintended consequences.

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