

From Regulation to Reality: Avoiding pitfalls in the EU AI Act

T. Chib, A.I. Hakkers and R. van Kempen¹

The EU AI Act promises to bring clarity and accountability to artificial intelligence, but its current framework leaves critical gaps that organisations must navigate with caution. How can companies ensure they meet these expectations without falling into the very gaps the regulation leaves open?

1. Introduction

With the EU AI Act (Regulation (EU) 2024/1689)² moving from legislative text to real-world enforcement, businesses across sectors are assessing how their AI systems fit into the evolving regulatory framework. The staged implementation means that obligations for general-purpose AI (GPAI) models³ have already been in effect since 2 August 2025, while high-risk medical AI systems will face compliance requirements from August 2026. This phased timeline⁴, combined with ongoing consultations and evolving guidelines, rewards early action. Companies that move now can influence emerging interpretations and establish defensible compliance positions through well-documented decisions. Those that wait face the downsides of conflicting guidance, rushed implementation, and precedents set by competitors.

In this article we converge three complementary perspectives, legal and governance, cybersecurity, and anonymisation to take a joint view on the EU AI Act, highlighting recent regulatory developments, persistent ambiguities, and practical risks that companies must address to comply without stifling innovation. We refer to this framework as the 'Golden Triangle' framework. This triangular relationship structures our analysis of six critical pitfalls organisations must navigate.

a. Legal/Governance Perspective

Examines regulatory requirements, compliance obligations, and governance structures needed for AI Act adherence.

b. Technical/Security Perspective

Considers system boundaries, technical documentation requirements, and security implications of compliance.

c. Anonymisation/Data Utility Perspective

Analyses how compliance requirements affect AI model performance, data processing capabilities, and innovation potential.

2. Grandfathering: A compliance loophole?

The EU AI Act's approach to existing high-risk AI systems through grandfathering provisions creates compliance complexities that surface during implementation. Article 111 of the AI Act stipulates⁵:

“...this Regulation shall apply to operators of high-risk AI systems, other than the systems referred to in paragraph 1 of this Article, that have been placed on the market or put into service before 2 August 2026, only if, as from that date, those systems are subject to significant changes in their designs...”

The provision implies a regulatory exemption whereby high-risk AI systems predating August 2, 2026, may operate outside Chapter 3 requirements, contingent upon maintaining their original intended purpose and architectural design. Recital 177 of the AI Act clarifies that “significant change”⁶ is equivalent to “substantial modification”⁷ as de-

1. Tanya Chib is a regulatory and privacy lawyer at Privacy Rules, reachable at tanyachib@proton.me, Dr. Anna Hakkers is a cybersecurity expert focusing on strategic data security, reachable at ai.hakkers@gmail.com, and Renate van Kempen is a data anonymisation consultant at Business AI Made Easy, focusing on health data protection and reidentification risk assessment, reachable at renate@baime.nl.
2. European Parliament and Council, 13 June 2024, Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), OJ L 2024/1689
3. European Commission, 'Guidelines for general-purpose AI providers on the AI Office template for downstream providers', Digital Strategy, 31 July 2025, <https://digital-strategy.ec.europa.eu/en/policies/guidelines-gpai-providers>.

4. 'EU AI Act Implementation Timeline', artificialintelligenceact.eu, 1 August 2024 <https://artificialintelligenceact.eu/implementation-timeline/>.
5. European Parliament and Council, 13 June 2024, Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), OJ L 2024/1689, Article 111
6. European Parliament and Council, 13 June 2024, Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), OJ L 2024/1689, Recital 177.
7. European Parliament and Council, 13 June 2024, Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), OJ L 2024/1689, Article 3(23).

fined in Article 3(23) of the AI Act, meaning a change not foreseen or planned in the initial conformity assessment.

2.1. Practical Impact

An AI-powered medical device, for example, could avoid the stricter high-risk obligations by remaining unchanged, even if its classification and potential risks remain high. This raises critical questions: Will companies rush through significant changes before the deadline, make only incremental updates to avoid triggering compliance, or delay needed updates entirely to sidestep obligations? Any of those scenarios risks undermining the Act's long-term accountability goals.

Another risk comes from legacy systems that remain "frozen" to avoid triggering substantial modification which can consequentially lock in outdated security architectures. Unpatched dependencies, end-of-life components, flat network designs and weak audit trails widen the attack surface while the system stays formally unchanged. In the life sciences and healthcare sector, for example, the pressure not to alter certified builds for medical devices can create a perverse incentive to defer security upgrades, even as exposure grows. The result is security debt that compounds over time and undermines post-market monitoring, incident response and vendor assurance across the supply chain.

3. Defining an "AI system"?

The EU AI Act's effectiveness hinges on a deceptively simple question: what constitutes an AI system?

Despite the European Commission's Guidelines on the Definition of an Artificial Intelligence System (CIE Guidelines)⁸, organisations face critical gaps in both definition and scope that create compliance uncertainty. These ambiguities manifest in three interconnected challenges: (i) determining which algorithms qualify as AI systems versus traditional software, (ii) establishing boundaries in complex multi-model architectures, and (iii) defining where an AI system begins and ends for compliance purposes.

The following analysis examines how these definitional gaps create practical compliance challenges and proposes risk-based approaches for navigating the uncertainty.

The question of what is in or out of scope remains one of the AI Act's biggest grey areas.

3.1. Algorithm Types

While the CIE Guidelines exclude systems that 'infer in a narrow manner' (para. 5.2) and 'simpler' rule-based systems (para. 48), paragraph 39 nonetheless acknowledges that AI employing predefined rules and logical inference can be covered.

This inconsistency makes it hard for companies to determine whether common tools, such as spam filters, antivirus software, or other machine learning-powered services, should be in their AI inventories. The CIE Guidelines tell us to assess a system's ability to infer, analyse, and adjust, but without clear thresholds, this could easily become a compliance guessing game.

3.2. Multi-model architectures

Modern AI often combines multiple models. The CIE Guidelines do not clarify whether a combination of models is one AI system, how to set boundaries when integrating models from different providers, or whether orchestration layers count.

For complex deployments involving multiple models, such as financial platforms combining large language models for document analysis, proprietary credit-scoring algorithms, and third-party fraud detection systems, organisations should assess each component individually while considering the integrated system's overall functionality. The controlling factor should be whether the combined system demonstrates inference capabilities and environmental influence that exceed the sum of its individual components.

Compliance teams should consider documenting system boundaries based on functional integration rather than technical architecture. Where orchestration layers coordinate multiple models to achieve unified business objectives, the entire integrated system can be typically assessed as a single AI system for regulatory purposes. Conversely, where individual models operate independently with minimal integration, separate assessments may be appropriate.

Technical teams should collaborate closely with compliance personnel to ensure assessments capture both functional capabilities and intended use cases. The AI Act⁹ provides that systems designed for research and development purposes may qualify for regulatory exclusions, making intended use documentation critical for accurate classification.

Organisations should implement periodic reassessment procedures, particularly for systems undergoing significant modifications or deployment changes. Substantial system modifications may al-

8. European Commission, 'Guidelines on the Definition of an Artificial Intelligence System under the AI Act', February 6, 2025, para. 5.2.

9. European Parliament and Council, 13 June 2024, Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act), OJ L 2024/1689, Article 2(8).

ter regulatory classification, requiring ongoing monitoring rather than one-time assessment. Finally, compliance teams should maintain comprehensive documentation supporting classification decisions, including the rationale for determinations and any technical assessments conducted during the evaluation process.

This systematic approach enables organisations to navigate AI Act classification requirements through structured assessment methodologies while maintaining appropriate documentation for regulatory compliance purposes.

3.3. Where does an AI system begin and end under the EU AI Act?

The question “Where does an AI system begin and end under the EU AI Act?” sounds simple. In practice, it shapes how organisations build, secure and govern their technology. Once a system falls within scope, compliance turns on the boundary that is drawn around it. Draw it too tightly and obligations are missed. Draw it too loosely, and teams are burdened with controls that do not fit the risk.

The scope matters for day-to-day operations. Compliance teams need to know which parts of the estate must meet the Act’s requirements, especially when the use case is high risk. Risk managers need clarity on what to assess and test. Technical writers need to document the right components so that auditors can follow the thread from design through deployment. Clear edges also support accountability. If responsibilities are ambiguous, gaps open up in maintenance, incident response and post-market monitoring. The same is true for the supply chain. Without a view of which third-party elements form part of the system, vendor assurance becomes guesswork. These decisions should also map to the Act’s operator roles, clarifying who is the provider, who is the deployer and where other responsibilities sit.

The CIE Guidelines on the definition of an AI system offer a helpful anchor. In essence, they include both hardware and software as part of an AI system. Processors, memory, storage, networking and input or output interfaces are in play because they make computation possible (para. 11). So are the software components that handle how the hardware processes data and performs tasks (para. 11). The guidelines require thinking across the lifecycle. The components necessary for training, validation, deployment, and operation can all be part of the same system even if they do not run at the same time (para. 10-12). They also tie scope to the system’s internal objective, which is the result the system is designed to achieve (para. 24-25). If a component is needed for the system to achieve what it is designed to do, it should be considered when the perimeter is defined.

Even with that steer, grey areas remain. Consider data infrastructure. A great deal of value is created before data ever reaches a model, through ingestion,

transformation, and feature engineering. Imagine a real-time fraud platform that draws on streaming pipelines, which extract features before inference. Are those pipelines within the AI system? The CIE Guidelines don’t clearly say whether data-processing infrastructure, beyond the direct compute, belongs in scope.

One way to address this is to use the following approach. If those steps merely prepare data and do not themselves carry out inference or learning, they often sit beside the AI system rather than within it. Where feature engineering is itself integral to achieving the internal objective, parts of that pipeline can fall inside scope. They still shape outcomes. A prudent operator will map and monitor the interaction and will explain why the system edge sits where it does.

Human interaction raises a similar puzzle. The hardware interface is easy to see. The software layer through which people consume outputs or send feedback can be harder to place. Take a customer service assistant that collects thumbs-up or thumbs-down ratings and uses them to guide future behaviour. The CIE Guidelines recognise input/output interfaces as hardware but leave unresolved whether the software interface through which users engage the system falls within the AI system boundary.

The following rule of thumb can be considered. If that interface is necessary to deliver outputs and to capture signals that shape performance, the most defensible view is that it sits within scope for compliance and documentation. Treating it as an external accessory understates its influence.

Adaptive operation complicates matters again. Many systems continue to learn after deployment, whether through scheduled retraining or incremental updates. Picture a bank’s fraud controls that fold new transaction patterns and analyst labels into a monthly refresh. From CIE Guidelines it’s unclear whether monitoring tools and feedback mechanisms should be considered part of the system.

This is how the system boundary could be defined. The monitoring tools that track drift, the workflows that gather feedback and the processes that approve and push new models are not decorative. They are how the system sustains its objective over time. They belong in scope for risk management, change control and technical documentation, even if they run on a different cadence to the inference service.

Cloud deployment does not dissolve the perimeter. If the system runs on serverless functions or containers, the compute and networking environment is part of the machine basis in which it operates. The shared responsibility model between provider

and customer changes who does what, not whether the component matters. The technical file should describe runtime characteristics, interfaces, control measures and the allocation of responsibilities.

The sensible approach is to document runtime characteristics and interfaces, then assign obligations contractually, so the controls match the architecture. Shared resources and multipurpose sensors are another grey zone. The fact that a CPU hosts both AI and non-AI workloads does not mean the entire host is the compliance perimeter. The same camera might support safety functions and unrelated analytics. What matters is the configuration that enables the AI system to achieve its objective, the data pathways it uses and the behaviours it can trigger. The obligations attach to the system operated and the parts of the environment that it relies upon. Configuration control should be maintained over shared components so that changes elsewhere do not alter risk without being noticed.

The definition of what constitutes an AI system from guidelines of the European Commission on AI systems provides important guidance, but organisations must still make practical judgments about where to draw system boundaries when it comes to grey areas.

What emerges from the examples given in this article is a practical way to set the perimeter. Start with purpose. Define the objective the system is meant to achieve and trace the components without which it cannot do so. Include the elements that deliver outputs to users or other systems, and the elements that sustain performance over time. Map the interfaces to upstream and downstream platforms and record why they have been placed in or out of scope. When third parties are relied on, allocate responsibilities with enough precision that an auditor could test them. This approach keeps the focus on outcomes and risk rather than on labels or organisational charts. Sector context matters, and in healthcare technology sector, financial services and critical infrastructure, domain rules may influence where the perimeter is set.

As AI Act implementation advances, expect clearer answers from rule makers, the courts, and industry practice. Until then, what matters is a clear, reasoned scope and the discipline to hold to it. Document judgment and explain how the parts fit together. Show that the control environment matches the way the system actually works. That is what good governance looks like, and it is the best defence when the rules meet reality.

In short, a risk-based, purpose-driven boundary should be set, which includes the components needed to achieve the system's internal objectives and to deliver and act on outputs, and document all external interfaces.

4. **GPAI Code of Practice: lessons for compliance strategy**

The development of the EU's General-Purpose AI Code of Practice offers a case study in how regulatory processes can be compromised, with implications for compliance requirements in the evolving AI governance landscape.

After a nine-month drafting process involving over 1,400 stakeholders, the European Commission's AI Office published guidance on July 10, 2025, that was weaker than earlier drafts due to exclusive access granted to leading US technology companies in the final stages¹⁰. This resulted in the removal of critical safeguards, including emergency preparedness requirements that are standard practice in other high-risk industries, measures for systemic risk mitigation during development, and explicit whistleblower protections. This demonstrates three key takeaways: first, the shift from rule-based to outcome-based frameworks, while offering flexibility, significantly raise the enforcement bar and requires organisations to document their compliance methodologies with exceptional rigor; second, the Code's problematic requirement that providers share full Model Reports only after deployment perpetuates a "deploy first, question later" mentality that increases both political and financial costs of corrective action; and third, successful compliance in this environment demands proactive engagement with the spirit, not just the letter, of regulations.

Rather than viewing standards as an opportunity for minimal compliance, forward-thinking organisations should anticipate that courts and institutions will likely rely on the Codes immediately to assess compliance, making them the de facto baseline for responsible AI behaviour globally. The most strategic approach involves implementing robust internal governance that exceeds current requirements, as regulatory standards inevitably strengthen over time, and early voluntary adoption of higher standards provides competitive advantage while reducing long-term compliance risk.

5. **Managing dual obligations for downstream providers**

The EU AI Act's implementation creates an unprecedented compliance landscape where downstream providers integrating GPAI models face a complex dual obligation structure that demands risk management and vendor oversight strategies.

The AI Act establishes a transparency framework requiring GPAI providers to create and actively provide documentation specifically designed for downstream AI system providers, including the model's

10. European Commission AI Office, 'General Purpose AI Code of Practice', July 10, 2025.

intended tasks, acceptable use policies, technical specifications, integration requirements, and information about training data provenance¹¹. However, this upstream compliance does not absolve downstream providers of their own obligations. Companies that merely use AI systems, especially in high-risk applications such as recruitment, healthcare, or critical infrastructure are required to maintain complete inventories of the systems they use and ensure prohibited applications are not deployed.

This dual structure creates a compliance chain where downstream providers must simultaneously verify upstream GPAI compliance while ensuring their own AI system meets applicable risk-based requirements. The challenge intensifies because GPAI obligations became enforceable on August 2, 2025, with full enforcement including fines beginning on August 2, 2026, while comprehensive high-risk AI system obligations follow one year later in August 2027.

5.1. Vendor Management and Due Diligence Imperatives

The timing disconnect between GPAI and AI system obligations creates practical challenges for compliance teams. Downstream providers integrating GPAI models into AI systems must cooperate with upstream providers to ensure compliance, requiring the GPAI provider to provide certain information to enable downstream compliance. This necessitates robust vendor management processes that go beyond traditional due diligence approaches.

Downstream providers must establish verification mechanisms to confirm that GPAI suppliers maintain appropriate upstream compliance, including adherence to transparency requirements, copyright compliance policies, and for systemic risk models safety and security measures.

Organisations conducting substantial modifications to existing GPAI models, including retraining, fine-tuning, or comparable technical interventions that materially alter system functionality, performance characteristics, or risk profiles assume provider status under the AI Act, thereby inheriting the complete obligation set applicable to original GPAI developers.

5.2. Strategic Compliance Recommendations

Given the regulatory complexity, compliance teams should implement several strategic measures. First, establish comprehensive vendor assessment protocols that verify GPAI provider compliance with Articles 53 and 55 obligations, including documentation

requirements, copyright policies, and systemic risk management measures where applicable. The European Commission has indicated that adherence to approved Codes of Practice will be favourably considered when assessing compliance and determining regulatory fines, making GPAI supplier adherence to the voluntary Code of Practice a critical evaluation criterion.

Second, develop internal governance frameworks that anticipate both immediate GPAI obligations and future high-risk AI system requirements. This includes implementing AI system inventories, risk categorisation processes, and documentation standards that will support compliance across both regulatory timelines. Deployers bear the obligation to implement suitable technical and organisational safeguards to ensure their utilisation of high-risk AI systems incorporating general-purpose AI models aligns with the downstream provider's specified usage protocols.

Third, establish clear contractual frameworks with GPAI suppliers that allocate compliance responsibilities and ensure access to necessary documentation and technical specifications. The AI Office may develop recommended standard contract terms applicable to high-risk AI system providers and related entities, offering voluntary adoption mechanisms that support compliance-driven contractual arrangement

6. GDPR and legitimate interest in AI

The intersection of the EU's General Data Protection Regulation ("GDPR") obligations and AI Act compliance creates a complex regulatory landscape that requires sophisticated risk management approaches. The European Data Protection Board's ("EDPB") Opinion 28/2024 on AI models provides critical guidance for organisations navigating these overlapping frameworks, particularly regarding the use of legitimate interest as a legal basis for AI development and deployment activities.

6.1. The EDPB's Three-Step Framework for Legitimate Interest Assessment

The EDPB's Opinion 28/2024, adopted on December 17, 2024, establishes a three-step assessment framework for controllers seeking to rely on legitimate interest under Article 6(1)(f) GDPR¹² for AI model development and deployment. This framework requires controllers to demonstrate: (1) pursuit of a legitimate interest that is lawful, clearly articulated, and real rather than speculative; (2) necessity of processing, meaning no less intrusive alternatives exist to

11. European Commission, 'Guidelines on the Scope of Obligations for Providers of General-Purpose AI models under the AI Act', July 18, 2025.

12. European Parliament and Council, 27 April 2016, Regulation (EU) 2016/679 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation), OJ L 119, Article 6(1)(f).

achieve the same purpose; and (3) successful completion of a balancing test ensuring the legitimate interest does not override data subjects' fundamental rights and freedoms.

The EDPB specifically recognises legitimate interests that may apply to AI contexts, including developing conversational agents to assist users, creating fraud detection systems, and improving cybersecurity threat detection capabilities. However, controllers must demonstrate that processing personal data is strictly necessary for these purposes and that no alternative approaches using synthetic or anonymised data could achieve the same objectives effectively.

6.2. Critical Compliance Challenges for Modern AI Development

The Opinion addresses fundamental challenges facing AI developers, particularly those utilising web scraping and large-scale data collection. The EDPB emphasises that reasonable expectations of data subjects play a crucial role in the balancing test, noting that the complexity of AI technologies makes it difficult for individuals to understand the variety of potential uses and processing activities involved. The assessment of reasonable expectations must consider whether personal data was publicly available, the nature of any relationship between data subjects and controllers, the context of data collection, the source from which data was collected including privacy settings offered, potential further uses of the model, and whether individuals are actually aware their personal data is online.

6.3. Cross-Regulatory Risk Management Strategies

The convergence of GDPR and AI Act obligations requires integrated compliance approaches that address both frameworks simultaneously. Controllers deploying AI models developed by third parties must conduct appropriate assessments to ascertain that models were not developed through unlawful data processing, particularly where infringements have been determined by supervisory authorities or courts. This creates a compliance chain requiring sophisticated vendor due diligence processes that verify upstream GDPR compliance alongside AI Act obligations.

When organisations rely on legitimate interest for AI deployment, the lawfulness of initial training data processing must be factored into the legitimate interest assessment, as unlawful development phase processing may impact the lawfulness of subsequent deployment activities. This necessitates comprehensive documentation of AI model provenance

and training methodologies, extending beyond traditional technical documentation to include detailed legal compliance records.

6.4. Practical Implementation Framework for Compliance Teams

Compliance professionals should implement several strategic measures to navigate these overlapping requirements effectively. First, establish comprehensive AI governance programs that integrate GDPR and AI Act compliance assessments from the earliest development stages. This includes conducting legitimate interest assessments that specifically address AI-related risks and implementing technical safeguards that exceed minimum legal requirements.

Second, develop robust vendor management protocols that verify both GDPR compliance in model training and AI Act adherence in model deployment. Organisations must document their assessment processes and maintain evidence that they have evaluated potential GDPR infringements in their AI supply chain, as supervisory authorities will scrutinise these accountability measures when investigating specific AI systems.

Third, implement enhanced transparency measures that go beyond standard GDPR information requirements. The EDPB recommends additional safeguards including public communications about data collection criteria, alternative forms of informing data subjects through media campaigns and transparency reports, and systematic presentation of information through model cards and transparency labels.

The integration of GDPR and AI Act compliance requirements represents both challenge and opportunity for organisations developing AI systems. By implementing governance frameworks that address both regulatory streams, maintaining detailed documentation of legal compliance measures, and adopting technical safeguards that exceed minimum requirements, organisations can navigate the current landscape while positioning themselves for success as enforcement practices continue to evolve across European markets.

7. Who decides if an AI model is anonymous?

Establishing when an AI model is truly anonymous is one of the most unsettled questions in the current regulatory debate. The EDPB Opinion¹³ 28/2024 makes clear that a model can only be considered anonymous if both the risk of extracting personal data and the likelihood of a successful attack are insignificant. Yet there is no recognised methodology,

13. European Data Protection Board, 'Opinion 28/2024 on certain data protection aspects related to the processing

of personal data in the context of AI models', 17 December 2024.

independent certification body, or standardised process for conducting such an assessment.

This creates practical uncertainty. Should companies commission external reviews, rely on internal privacy teams, or engage independent technical auditors? And which authority ultimately validates those decisions, national data protection authorities, the AI Office, or sector regulators?

Adding to the challenge is the mosaic effect¹⁴: even when datasets are anonymised, combining them with external sources can re-identify individuals. This risk is especially acute for AI models trained on rich, multi-source health or behavioural data. Furthermore, anonymisation is context-dependent¹⁵; what is safe within a secure environment may no longer be safe once data is used to train models or shared across borders.

Research is also moving fast. New privacy-preserving techniques are emerging, such as penalty-driven or utility-driven anonymisation, which try to optimise the trade-off between privacy and model performance. While promising, these approaches are not yet standardised or tested at regulatory scale, leaving organisations uncertain about whether they meet the high GDPR threshold of negligible risk.

Why it matters: Until clear methodologies and oversight mechanisms are established, businesses should treat anonymity assessments as a high-risk compliance area. Documenting assumptions, engaging multidisciplinary experts, and adopting conservative interpretations of “anonymous” are the best defences against regulatory scrutiny and against creating models that are neither fully private nor fully useful.

The AI Act draws the map, but organisations must still chart the route.

8. Final Words

The EU AI Act is a milestone in AI governance, but its effectiveness will depend on how companies bridge the space between regulatory intent and operational reality. The six critical gaps we have examined demonstrate a complex compliance landscape. Seen through our Golden Triangle lens, the sharpest risks surface where legal governance, technical architecture, and data utility meet: the grandfathering of high-risk systems such as medical devices, the CIE Guidelines’ grey areas on system definition and multi-model orchestration, the GPAI Code of Practice’s outcome-based tilt and post-deployment model reporting, the downstream burden to verify upstream claims while meeting one’s own duties, and the EDPB’s legitimate-interest tests and unsettled standards for model anonymity.

Priority next steps for organisations are to build and maintain an AI inventory with purpose-driven, functional boundaries and the rationale for each inclusion or exclusion; set periodic reassessments so “significant changes” are managed deliberately rather than avoided; require GPAI suppliers to provide usable documentation and contract terms that are proportionate to the risks and complexity of their AI models, then test what has been given; integrate GDPR and AI Act workflows by running legitimate-interest assessments alongside provenance logs, model cards, and clear transparency notices; treat anonymisation as a high-risk area with conservative thresholds and multidisciplinary review; include monitoring, feedback, and retraining processes in scope and in the technical file; and adopt internal standards that exceed today’s regulatory requirements so one is well positioned in the evolving AI landscape.

The organisations that will succeed are those that act early, document rigorously, and embrace a risk-based approach that goes beyond minimum compliance. In short, the Act sets the direction, but it is the choices made now, in governance, data protection, and technical design, that will determine whether AI in Europe becomes both innovative and trustworthy.

14. Gary LaFever, ‘Beyond GDPR: Unauthorized reidentification and the mosaic effect in the EU AI Act’, International Association of Privacy Professionals (IAPP), 15 March 2024, <https://iapp.org/news/a/beyond-gdpr-unauthoriz>

15. Santa Borel, ‘How Context Affects Anonymization in AI Model Development’, January 13, 2025, <https://privacy-analytics.com/resources/articles/how-context-affects-anonymization-in-ai-model-development/>