



Domino



Addressing the EU AI Act:

Strategies for compliance and accelerated business value



Table of contents

Executive summary	03
Introduction	03
What is the EU AI Act?	04
EU AI Act compliance simplified with Domino	06
More regulations to come	10
Conclusion	11



Executive summary

This paper outlines a set of key considerations for IT, risk, and line of business executives tasked with addressing the EU Artificial Intelligence Act (AI Act), and driving business transformation with AI while de-risking AI. The report demonstrates how Domino Data Lab's Enterprise AI Platform helps global enterprises ensure AI Act compliance and quickly fosters responsible and safer AI.

Introduction

On March 13, 2024, the EU voted for the landmark [EU AI Act](#), which establishes a robust regulatory framework for mitigating AI risks and implementing governance. This act will impact how multinational enterprises navigate the process, people, and technologies used to build, operate, and govern thousands of predictive AI and generate AI (GenAI) systems and models applications globally.

The EU AI Act establishes a legal framework for AI systems in the European Union, focusing on risk management, transparency, data governance, and ethical AI practices. As AI technologies evolve, the need for robust solutions to ensure compliance and foster responsible AI becomes imperative. Domino's [Enterprise AI Platform](#) offers a range of features that address these needs, and go beyond them to establish safe and responsible AI in the enterprise.

All global enterprises must now navigate the complexities of complying with the EU AI Act if they do business in EU countries.

With hundreds of highly-regulated customers in the pharmaceutical, financial services, and government sectors deploying and managing AI in production in Domino, the Domino team understands what it takes to govern AI safely and at scale.

Beyond addressing individual regulations like the EU AI Act, enterprises must establish responsible AI systems as a part of their long-term risk mitigation strategy. Companies improve the business impact of AI and de-throttle innovation by adopting technologies that are open and extensible, to support both emerging regulations and their own evolving corporate risk policies. By thinking more broadly about putting responsible AI policies into practice throughout the end-to-end AI lifecycle, companies can not only comply with regulations today and tomorrow, but they can also de-risk AI and generate more long-term business value from it.

Keep reading.
We'll outline key
strategies for
addressing the
the EU AI Act






What is the EU AI Act?

In March 2024, the European Union (EU) passed [artificial intelligence \(AI\) legislation](#), called the EU AI Act. These are considered the world's most sweeping AI regulations, with enforcement beginning in late 2024 into early 2025. The AI Act aims to mitigate the negative impacts of rapidly deployed high-risk AI systems and applications, without throttling innovation. This landmark regulation could create a de facto standard for other jurisdictions to follow, by providing a responsible AI framework that mitigates negative impacts such as bias, data leakage, and other risks posed by standalone and composite AI systems (a combination of predictive AI and GenAI systems like large language models such as ChatGPT).

The EU AI Act categorizes AI systems and applications into three categories, based on risks posed: 1) high-risk AI, 2) limited risk, and 3) minimal risk.

High-risk AI systems include safety components of previously regulated products and stand-alone AI systems in specific areas (e.g., critical infrastructure), which could negatively affect or significantly harm the health, safety, and fundamental human rights of people, or the environment. Limited risk AI systems include those used to manipulate or deceive, such as chatbots, or GenAI systems that produce false content such as deep fake imagery. Lastly, minimal risk systems are those systems posing limited threats, such as spam filters — these have no mandatory restrictions or obligations, under the EU AI Act.

EU AI Act: risk levels



What is the EU AI Act?

The EU AI Act imposes obligations on users of high-risk systems, impacting how these companies develop, monitor, and maintain their AI models, systems, and applications. In addition, the AI Act outlines steep and highly enforceable penalties, including fines of 1.5% and up to 7% of global turnover (total revenues). Stricter regulations are placed on high-risk systems and applications, and the law allows an outright ban of those that pose unacceptable risks. Low-risk applications and those using open source software will face fewer requirements, for example (e.g., video games and spam filter applications).

The AI Act applies to providers of “high-risk” AI systems and applications, including GenAI products like Open AI’s ChatGPT and DALL·E and [Google’s](#) Gemini. These products must be subjected to tests that detect bias and inaccuracies and ensure transparency before companies use these tools to put products into market.

The AI Act imposes obligations on those who develop and deploy AI models; both parties are subject to requirements ensuring the safe, secure, and transparent operation of AI. Under the AI Act, these AI systems come with mandated obligations such as adequate risk assessment and mitigation, use of high-quality datasets feeding the systems to minimize bias and discriminatory outcomes, logging of activity to ensure traceability of results, detailed documentation for compliance assessments, human oversight, and robust security and accuracy.

Beyond addressing EU AI Act compliance, Domino’s platform is continually monitored and audited against standards like SOC 2 Type 2 and ISO 27001:2013, and Domino supports GDPR and HIPAA compliance.



AI Act high-risk systems and applications:

Systems impacting lives, livelihoods, health, and more —



Credit and finances



Administration and justice



Employment



Insurance



Education and vocational training



Healthcare



Critical infrastructure and public services




EU AI Act compliance simplified with Domino

Domino unleashes the potential and power of AI for sophisticated and highly-regulated enterprises by making it safer, and less risky. Domino's unified platform helps executives orchestrate and scale AI, GenAI, and data science, by building, deploying, governing, and managing AI, GenAI, and data science workloads that deliver responsible AI by default.

Data science and research teams can access their preferred data, tools, computing resources, AI models, and projects across any environment. This allows them to foster collaboration, establish best practices, and track models in production while ensuring governance and reducing risks and costs. Additionally, by democratizing the essential tools and methods data scientists and analysts need to develop trustworthy AI and ML models (e.g., methods of explainability, detecting bias, mitigating bias, evaluating fairness, monitoring model drift and inaccuracies), Domino makes them more easily accessible, across large organizations.

Domino is the end-to-end enterprise AI platform for compliance:



Domino is designed to make AI responsible by default. Domino makes it easy to de-risk AI and govern AI workflows to address the EU AI Act by delivering an audit-ready platform with best-in-class reproducibility, turnkey model governance, monitoring, and remediation, and enforced compliance through robust security controls.

Enterprises have on flexible, system of record where they can quickly embed and enforce responsible and safer AI policies of the EU AI Act, across all workflows into every AI model produced — in any phase of the AI lifecycle (e.g., build, deploy, operate, fine-tune, maintain). And Domino is an open platform so enterprises can stay future-ready with new and emerging tools, with costly IT integration projects.

With Domino, global enterprises address not only current regulations like the EU AI Act, but also emerging regulations focused on risk detection, assessment, and mitigation. By proactively de-risking AI, enterprises de-throttle AI, thus accelerating time-to-business value and time-to-market for AI solutions.

Domino ensures that AI is responsible by default



Make AI audit-ready

All models are audit-ready with guaranteed reproducibility and embedded audit trails that record all lineage (code, data, environments, results, versions).



Enforce AI safety policies

Enforce AI policies with out-of-the-box templates that automate workflows for building and operating models safely and reliably.



Maintain consistent standards

Integrated model review, staging, testing and approval workflows help embed and enforce requirements for every model, every time.



Ensure fairness and accuracy

Model monitoring and remediation continually score model accuracy and detect drift. These mechanisms help prevent unfair and biased models from entering production environments.



“For the past ten years, Domino has been working with highly regulated industries and government customers like the U.S. Department of Defense and the Financial Industry Regulatory Authority (FINRA), a regulatory body concerned with financial risk. While many customers come to Domino to address new regulations, they stay because it makes good business sense to de-risk AI and establish safe best practices that accelerate AI’s business value.”

– Kjell Carlsson,
Former Forrester Analyst | Head of
AI Strategy, Domino



Here is how Domino addresses the EU AI Act's requirements for high-risk use cases:

1. Risk management system

Model registry | reproducibility engine | model monitoring

Assessing AI risks and mitigating them involves people, processes, and policies — beyond code and algorithms. Domino integrates and embeds your policies addressing EU Act requirements into AI model review, staging, testing, and approval workflows, making the regulation actionable and enforceable.

Domino offers a suite of capabilities for managing AI and ML model risks. For example, Domino Model Registry is a central registry for tracking all AI experiments, artifacts, and models, with built-in version control, so all AI is reproducible for audits and compliance. Data science, risk, IT and line of business leaders can view all models in development and production from a single pane of glass, regardless of where they were developed or hosted (on Domino or elsewhere). Lastly, Domino mitigates risk with integrated model monitoring, to detect model drift and accuracy (risk areas), and model approvals that enable promotion of only the best, least risky models to production.

2. Data governance

Centralized data access | role-based access control | lineage tracking

Large enterprises are unable to ensure data governance and control data access or provide accurate data lineage, across siloed data, data users, and data providers. Domino ensures data governance by centralizing access, control, and visibility across all datasets — structured or unstructured, stored in any infrastructure environment. Users can continue to query all data from one platform while minimizing data movement and

reducing risks. Domino makes all types of data sources, stored anywhere, more accessible without compromising safety or data privacy, with fine-grained, role-based access controls and encryption.

Lastly, Domino automatically tracks data lineage (versions, where it came from, how it changed), metadata, data provenance, and data access, and ensures security and compliance. Data audit trails automatically record all access to ensure model approvers can confidently ensure compliance with enterprise security and audit policies.

3. Quality management system

Model cards | model review and approval | model monitoring
| model staging

With Domino, data scientists have access to all the methods they need to evaluate, validate, and test their models, while collaborating with other users for independent validation, and continuously monitoring models to ensure quality.

In Domino, users can build, manage, and reproduce high-quality models to deliver high-performance AI. With a simple point-and-click interface, users can visualize, prepare, and transform data, and every step in the data preparation process comes with a detailed audit trail.

In terms of model quality, Domino tracks all model lineage and artifacts for guaranteed reproducibility. Model cards record the purpose, performance, limitations, and biases of all models,

automatically populating with critical details. Then, Domino Reviews require model assessments, and Domino Model Monitoring workflows ensure every model is assessed for accuracy, bias, and drift, and a score is assigned. Next, Domino Model Approvals require approvals and risk assessments before risky models are ever put into production. Lastly, Domino Model Staging ensures models are isolated during testing, to mitigate risks and the possibility of legal, reputational, and societal harm.

4. Human oversight

Integrated human-in-the-loop | model review and approval
| model monitoring

Domino Model Reviews can assign human reviewers (human-in-the-loop) as part of Domino Model Reviews and Domino Model Approvals before a high-risk model is inadvertently released into production. This not only creates accountability, but also enforces compliance with your human review policies since no model is released without a responsible individual's sign-off, recorded in Domino.

In Domino, managers also have visibility into gates passed for each model, from those being developed to those in production. In addition, Domino Model Monitoring notifies managers and model owners of model performance.



5. Technical documentation

[Model cards](#) | [reproducibility engine](#) | [traceability](#)

Domino Model Cards automatically capture critical details in a single view, including inputs, outputs, model specifications and performance, version and reviewer notes, and more so that models meet the EU AI Act requirements codified by an enterprise's policies. Each card evaluates a model's accuracy, fairness, compliance, and auditability. Reviewers can easily test a model via embedded model APIs, explore the experiment details, or open a workspace to dive deeper into the training context. This helps enterprises accelerate model reviews and approvals, and deploy models with confidence and complete, audit-ready documentation.

6. Monitoring

[Model monitoring](#) | [traceability](#)

Domino offers integrated model monitoring to help diverse, global teams track model quality, drift, and performance. Each model is tested using preferred tools, predictions are captured, and a score is assigned. Based on that score, models can be approved for promotion into production or rejected and sent to human reviewers for further validation or fine-tuning. In addition, Domino Audit Logging makes all user actions traceable.

7. Registration

[Model registry](#) | [model cards](#)

Domino Model Registry gives teams a unified view of all of their models. Teams and managers can view all models with their status, metadata, and evolution (e.g., code, data, parameters, inputs, outputs, model specs, environments, results, and versions) for auditability and tracking. Capture critical model

details. Domino Model Registry makes it impossible to lose track of models, whether they were developed in Domino or not. This helps teams streamline model reviews so models are not promoted to production before they are reviewed. It also simplifies documentation handoffs to auditors.

8. Info to downstream users for transparency

[Model registry](#) | [model cards](#) | [reproducibility engine](#)

Domino guarantees the reproducibility of all AI model results and experiments by automatically tracking all model artifacts, data, code, with detailed lineage and provenance information, making it easy for downstream users and subcontractors to understand the purpose, intent, goals, and performance of a model. This helps create transparency so that decisions based on models can be understood and traced throughout the value chain. This is particularly important for high-risk AI use cases that directly impact consumers and citizens, since model records and reproducibility may be requested by these parties, or required in future legal inquiries.

9. Robustness, accuracy, cybersecurity

[Unified security](#) | [model monitoring](#) | [role-based access control](#)
| [model staging](#)

Domino is designed for the most highly regulated industries, including regulatory bodies. It is designed with enterprise-grade security and continuous model monitoring that assures performance and reliability.

First, Domino reduces the threat surface by unifying all data and AI model access, for granular visibility and traceability. Then, Domino offers continuous model monitoring, so all models

are evaluated and scored based on conformance with AI Act requirements. Then, all models are recorded with lineage details that make them reproducible, for auditors. Finally, all models are tested in isolated environments during model staging, before they are placed into production, where risks are heightened.

Domino also offers robust, enterprise-level, end-to-end security with RBAC for data and model access and encryption at rest and in transit to protect data privacy. Domino also seamlessly integrates with existing hybrid cloud, multi-cloud, and multi-region infrastructure to avoid data movements that can compromise data locality and cause data security and compliance breaches.


Since GenAI models and large language models (LLMs) like ChatGPT pose greater risks to responsible and safe AI and AI Act compliance. Domino offers seamless, secure, controlled access to any hosted LLM with the ability to govern LLM endpoint access, token usage, and API keys so data leaks are prevented. In addition, all endpoint activities are automatically logged for audibility.

10. Conformity assessment

[End-to-end platform](#) | [reproducibility engine](#)

The AI Act defines conformity assessments that verify that high-risk systems comply with the AI Act requirements. Domino simplifies this by automatically recording model goals, performance, and all artifacts so all AI is reproducible and auditable. For example, this means that auditors can restore one part of a project to an old state, run current code on an older version of the data, and view and revert file revisions.

4 steps of EU AI Act compliance



Source: [EU AI Act website](#) and [detailed regulations](#)

More regulations to come



The Domino Enterprise AI Platform supports compliance with the EU AI Act, so enterprises can avoid steep fines for non-compliance (up to 7% of total global revenues/turnover). However, Domino also goes beyond supporting current regulations, to support new regulations with an end-to-end platform that mitigates a variety of AI risks.

Domino's open, extensible, future-ready platform is designed to seamlessly adapt new regulatory requirements and each company's own responsible AI policies into all AI models, data, tools, and processes. New AI regulations are ever-evolving and arising in many forms. For example, voluntary compliance in G7 countries may soon mean that all enterprises will face mounting pressures from formal rules, nation-specific regulations, and self-imposed requirements in the near future.

Already, the [AI Pact](#) calls for voluntary commitments to the EU AI Act before it is enforceable, and the U.S. White House has been pressuring key industries, such as healthcare, to voluntarily comply with the [White House Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence](#) released in 2023. In addition, GenAI has faced pressure from the OECD and G7 leadership to self-regulate developers of GenAI.


These new regulations and emerging self-imposed requirements carry common obligations across five key areas directly addressed by Domino:

- Data and model security and privacy protection are required at all stages of the AI lifecycle.
- Pre-release risk assessment, planning, and mitigation focused on training data and implementing guardrails — addressing bias, inaccuracy, and other potential harm.
- Documentation required at release, covering steps taken in development and regarding the nature of the AI model or system (capabilities, limitations, description of training data, risks, mitigation steps taken, etc.).
- Post-release monitoring and ongoing risk mitigation, focused on preventing inaccurate or other harmful generated output, avoiding discrimination against protected groups, and ensuring users realize they are dealing with AI.
- Minimizing environmental impact from energy used to train and run large models.

More regulations to come (cont.)

Domino addresses today's and tomorrow's AI regulations, with a strategic, long-term approach to risk mitigation. Domino's Enterprise AI Platform helps you integrate your responsible AI practices and policies into the end-to-end AI lifecycle — from building, operating, and maintaining AI models and systems. This not only makes your AI initiatives compliant with the EU AI Act, but it also makes good business sense by de-risking AI, improving governance and transparency, and shortening time to value so you can unlock the full potential of AI across your enterprise.

The real ROI of responsible AI



Conclusion

Domino Data Lab's comprehensive approach to AI development eases and accelerates compliance with the EU AI Act's requirements. However, Domino goes far beyond the AI Act compliance; it offers an integrated solution for organizations seeking to adhere to new and emerging regulatory standards, responsible AI best practices that are good for business, and best-in-class risk mitigation measures that de-throttle AI and accelerate its business value.

For more details on this topic, visit [Domino's Responsible AI webpage](#) to learn more →



About Domino Data Lab

Domino Data Lab empowers the largest AI-driven enterprises to build and operate AI at scale. Domino's Enterprise AI Platform unifies the flexibility AI teams want with the visibility and control the enterprise requires. Domino enables a repeatable and agile ML lifecycle for faster, responsible AI impact with lower costs. With Domino, global enterprises can develop better medicines, grow more productive crops, develop more competitive products, and more. Founded in 2013, Domino is backed by Sequoia Capital, Coatue Management, NVIDIA, Snowflake, and other leading investors.

Learn more at www.domino.ai →