Bayerischer
KI-Innovationsbeschleuniger

Risky Business or Safe Bets?

Recommendations for accelerated compliance with the EU AI Act based on the Risk Classification of 628 German AI Startups

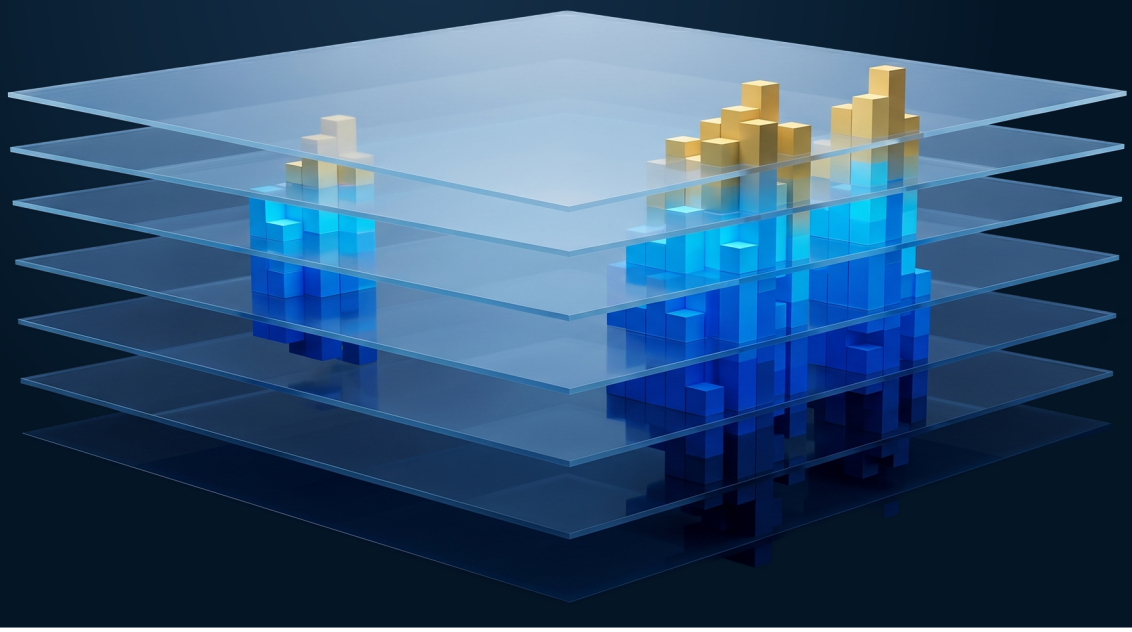

appliedAI institute for europe

Gefördert durch
Bayerisches Staatsministerium für Digitales

Content

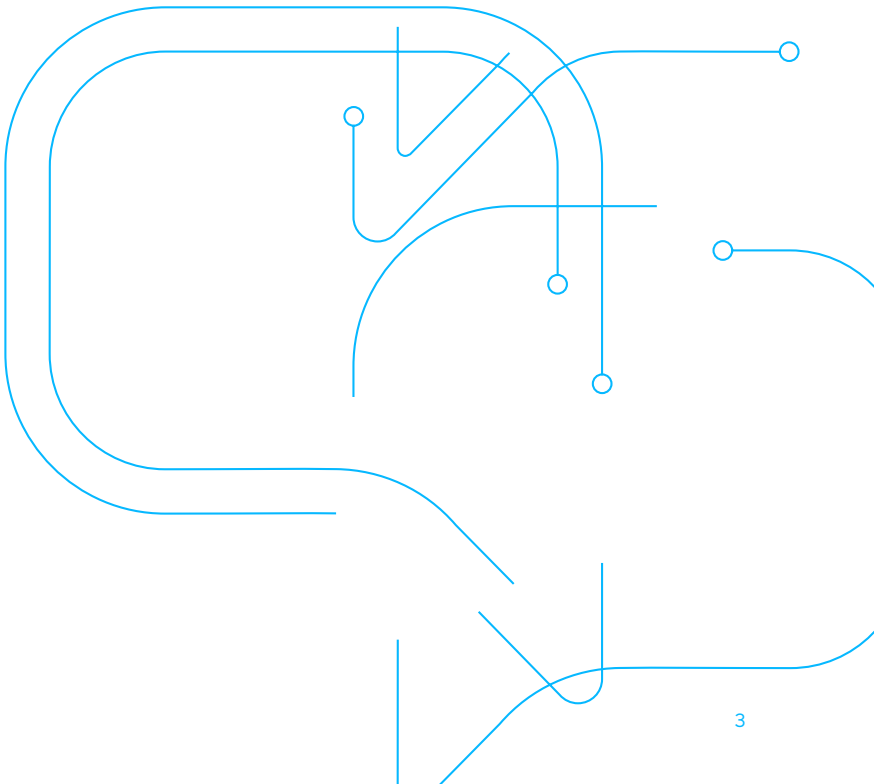

Foreword

With the EU AI Act now in force as of August 1st, 2024, European AI startups face a critical challenge: meeting extensive compliance obligations while navigating resource constraints and competitive pressures. This regulatory framework, while essential for ensuring Trustworthy AI development, presents a significant burden for startups already grappling with limited financial and human resources.

At the appliedAI Institute for Europe, we recognize the vital role AI startups play in driving Europe's innovation agenda. However, the added regulatory complexity introduced by the AI Act could hinder their growth and scalability. This study provides a comprehensive, data-driven analysis of the regulatory requirements these startups will face, offering the insights needed to design effective support mechanisms.

By analysing the risk classification of the entire German AI startup landscape imposed by the AI Act, we aim to bridge the gap in existing research and help inform strategic decisions that balance regulatory compliance with innovation. This report serves as a step towards ensuring that the EU remains a global leader in AI, while fostering an environment where startups can thrive.

This study was conducted as part of the Bavarian AI Act Accelerator, which is funded by the Bavarian State Ministry of Digital Affairs to support SMEs, start-ups, and the public sector in Bavaria in complying with the EU AI Act. As such, this study is an important contribution to the ongoing conversation between these actors, guiding us toward a future where innovation and compliance coexist.

Sincerely,

Dr. Wolfgang Fischer
Director Operations, Projects & Policy
appliedAI Institute for Europe

Executive Summary

The EU AI Act introduces a tiered regulatory regime for low-risk, transparency-obligated, high-risk, and prohibited AI systems; however, the distribution of AI use cases in Germany across these risk tiers remains unclear.

This study, which was conducted as part of the Bavarian AI Act Accelerator and funded by the Bavarian State Ministry of Digital Affairs, uses a novel LLM-based approach to classify the highest risk use case of 628 German AI startups under the EU AI Act and identifies where the regulatory burden is likely to concentrate across risk tiers, industry sectors and enterprise functions, federal regions, and funding profiles.

Since the AI Act entered into force, this study is the largest effort of its kind to classify AI systems. It is intended as a preliminary study to encourage further research into the regulatory implications of the AI Act and associated compliance costs for Germany's innovation ecosystem.

Key findings

- **The number of high-risk startups is higher than commonly assumed:** After excluding cases that could not be conclusively classified (69 out of 628, i.e. 559), nearly half of startups are classified as low-risk (48.3% = 270/559), while approximately one quarter are subject to transparency obligations (26.3% = 147/559) and approximately one quarter qualify as high-risk (25.4% = 142/559). No examined use cases were classified as prohibited.
- **Risk classification is frequently context-dependent:** 69 startups could not be conclusively classified due to incomplete information or because classification depends on deployment context and technical specifics that are not available in public descriptions.
- **Transparency obligations are likely to become a new baseline:** The study analyzes a subset of 111 generative AI startups and finds that 59% of them have transparency obligations and 12% have high-risk obligations, suggesting transparency requirements may emerge as a dominant "baseline" tier as LLM-based solutions proliferate and that the interplay between the Codes of Practice (CoP) for GPAI models and the requirements for high-risk AI systems will significantly influence downstream compliance.
- **Compliance burden concentrates in specific sectors and enterprise functions:** High-risk startups concentrate disproportionately in sectors like Human Health & Social Work Activities and Transportation, as well as in enterprise functions such as Human Resources and Operations. This concentration implies that compliance support measures tailored to sector and enterprise function are likely to be more impactful.
- **Regional ecosystem size does not materially change the overall risk-tier composition, but sectoral mixes differ and drive different governance needs:** Leading startup states show broadly similar proportional risk profiles (e.g., Berlin and Bavaria), yet sectoral composition differs by state, implying different supervisory expertise requirements.
- **Current funding does not yet reflect compliance costs:** Despite a quarter of all systems likely being high-risk, private funding patterns do not yet show any differentiation between risk tiers. However, this is a purely descriptive finding and does not necessarily reflect a lack of investor willingness.

Priority Recommendations

Stakeholder	Recommendations
EU / AI Office	• **Update assumptions and implementation support:** revisit impact assumptions on the share of high-risk systems and consider corresponding adjustments to supportive measures, including funding mechanisms. • **Delay high-risk and transparency obligations:** Delay the enforcement of the high-risk and transparency obligations as proposed in the Omnibus package so companies have the resources and time to implement the technical standards and Code of Practice for transparent systems.
German Federal Government / Federal Regions	• **Build supervisory capacity where burden concentrates:** equip competent authorities and market surveillance bodies with sufficient technical and legal expertise to handle a higher-than-expected volume of regulated systems. • **Accelerate practical guidance:** equip market surveillance bodies with a mandate to foster AI innovation; acknowledge best efforts by operators in good faith and with guidance. Penalties and fines should be the last resort. • **Coordinate across states on sector-specific use cases:** establish cross-state mechanisms to share expertise and converge on interpretations for the most affected sectors and enterprise functions.
Industry/ Startups and Investors	• **Treat classification as a lifecycle process:** AI startups should perform risk classification early in their product journey and be aware that product changes and deployment contexts can shift risk tier, particularly because many startups operate across various sectors and enterprise functions. • **Price regulatory exposure and compliance readiness into financing:** Investors should incorporate compliance costs and readiness signals into diligence and valuation.

Core limitations

- **Novel methodological approach:** This classification combines automated LLM analysis and human review. While this enables scalability, both have a margin of error.
- **Absence of guidance and incorrect classifications:** Interpreting the AI Act remains an ambiguous exercise until further guidance from the commission, particularly on identifying high-risk use cases, is published. It is possible that human reviewers incorrectly classified AI use cases.

1. Motivation

Prior to introducing the EU AI Act, the European Commission expected in its 2021 Policy Impact Assessment that most AI use cases will fall into the low-risk category, whereas “only 5% to 15% of all AI applications are estimated to constitute a high risk” (European Commission, 2021, p. 69)[1]. Since the publication of the AI Act, however, there have been no systematic efforts to assess or validate this expectation[2].

A higher than expected proportion of high-risk or transparency-obligated systems would imply greater compliance costs for European companies. This is particularly true for startups, which will likely be disproportionately affected by the regulatory and compliance obligations introduced under the EU AI Act.

For instance, a recent OECD report on Germany’s AI readiness has highlighted that German AI Startups are at a structural disadvantage in the global market with less access to venture capital, technical talent and computing power, especially in comparison to the United States (OECD, 2024). Compliance tasks will now require additional resources that AI startups are already struggling to gather. As of today, one can also observe that founders of generative AI Startups highlight the negative implications of the EU AI Act regarding additional resources and high regulatory complexity (Hutchinson *et al.*, 2024).

Our study aims to address this gap by providing a structured and data-driven assessment of the risk class distribution of German AI startups under the AI Act. We combine this assessment with data about these startups’ sectoral, enterprise, and geographic operations in Germany to better understand the regulatory consequences of the AI Act for these stakeholders[3].

We hope that empirical insights into the distribution of risk classifications across Germany’s AI startup ecosystem not only enhance transparency but also lays the groundwork for further research, more informed data-driven policy recommendations, and strategic interventions that can sustain Germany’s position as a leader in AI-driven innovation in line with European values.

1 See Ch. 2 for more details on risk classification.

2 Other studies conducted risk classifications based on prior versions of the EU AI Act, which significantly differs from the current legislation in terms of the criteria and rules for risk classification. For further insights regarding the analysis of prior versions of the EU AI Act, see Hauer *et al.* (2023) and appliedAI (2023).

3 This study was finalised before the European Commission's draft guidelines on high-risk classification under Article 6 (published 19 May 2026) and before the provisional agreement on the Digital Omnibus on AI (7 May 2026), which defers high-risk obligations to 2 December 2027 (stand-alone) respectively 2 August 2028 (embedded systems). Our classification reflects the framework available at the time of analysis; as the Omnibus adjusts timelines rather than the substantive classification rules, we do not expect a material effect on the reported distribution of risk classes.

2. Why the AI Act Risk Classification Matters for AI Startups

The obligations for an operator of AI systems under the AI Act depend on two key factors: the risk class of its system and its role in relation to that system.

In the general scheme of the AI Act, providers of AI systems - i.e, those that build and market them - face significantly higher obligations than deployers - i.e, those that merely use them. Apart from determining their role, operators must also classify their AI system into the AI Act's four-tiered classification framework. The goal behind this framework is to ensure that obligations are proportionate to the risk - the higher the risk of a given AI application, the greater the need to build it safely.

Our observations support the view that the majority of AI startups are likely to be providers. If a higher-than-expected share of them are building high-risk AI applications, the ecosystem as a whole will confront higher compliance costs. Accordingly, this chapter outlines the general risk classification framework under the AI Act, examines its financial consequences, and analyzes what this means specifically for AI startups in Germany.

2.1 General Risk Classification under the AI Act

Generally, under the EU AI Act all operators of AI systems are required to classify their systems according to a four-tiered risk framework. This classification is based on the level of risk an AI system poses to health, safety, and fundamental rights, as determined by its intended purpose.

The framework establishes the following categories:

Unacceptable-risk

Unacceptable-risk AI systems are prohibited under article 5 of the AI Act. These are systems that pose a clear threat to individuals or society, such as those that manipulate behavior through subliminal techniques, exploit vulnerable populations, or facilitate government-run social scoring.

High-risk

High-risk AI systems are permitted but subject to strict regulatory requirements. There are two routes to this classification under article 6 of the AI Act: systems that are part of products regulated under existing EU product safety laws (listed in **Annex I**), and standalone systems used in sensitive areas outlined in **Annex III**, such as education, employment, law enforcement, and critical infrastructure.

Transparency Obligations

Certain AI systems face specific **transparency obligations** under article 50 of the AI Act. For instance, users must be clearly informed when they are interacting with AI, such as chatbots, and synthetic content generated must be watermarked.

Low-risk

Low risk AI systems, i.e all systems which do not fall into any of the previous three categories, are largely unregulated under the AI Act. While no mandatory requirements apply, providers are encouraged to adopt voluntary best practices under article 95 of the AI Act.

2.2 The Costs of Being High-risk

Every piece of legislation comes with a cost. Startups that develop high-risk AI systems are subject to a broader range of stringent requirements compared to their lower-risk counterparts. Studies that supported the European Commission's impact assessment in 2021 estimated that "...for an enterprise of 50 employees, the total cost for setting up and maintaining compliance with the AI Act ranges from EUR 216,000 to EUR 319,000 for one product in year one." (Renda *et al.*, 2021, pp. 152–153).

While estimating the cost of AI Act compliance is difficult (Laurer *et al.*, 2021; Mueller, 2021) — and highly dependent on the conformity assessment process — startups often face high costs for structural reasons. For one thing, compliance efforts cannot often be addressed in isolation and may require external consultants or lawyers. Additionally, startups that did not typically operate in regulated environments might confront a far steeper learning curve, requiring upskilling and additional manpower to build new features necessary for compliance. Finally, startups might not have the know-how or the financial resources to navigate engagement with notified bodies and the certification process. Startups will confront these upfront and ongoing costs at the same time that they are heavily dependent on external funding to validate and scale their business models.

2.3 Support for AI Startups

While all operators of AI systems must meet the same obligations under the AI Act, there are special provisions that tailor to **SMEs and startups**. Recognising that these stakeholders have an important role to play in innovation, the Act requires EU member states, the European Commission and other bodies to provide legal and technical support to them.

These provisions can broadly be classified into four types (see chapter VI of the AI Act for more details):

- **Information dissemination:** EU member states are required to organise tailored training activities for SMEs and startups and to set up dedicated channels to respond to queries about the law and its implementation.
- **Reducing compliance costs:** SMEs and startups will benefit from a lower standard of rigour for technical documentation and from lower fees from relevant authorities during the compliance process.
- **Participation in decision making:** Member states and the EU Commission are required to facilitate the participation of SMEs and startups in decision making fora, like the advisory forum and standard setting institutions.
- **Priority access to regulatory sandboxes:** Startups and SMEs will benefit from easier access to regulatory sandboxes—which are controlled environments where businesses can experiment with products whose regulatory status is uncertain[4].

While most of the AI Act will become applicable in August 2026, there is no single 'go live' date for all of these measures. It is also unclear to what extent these measures will prove cost-saving for startups because there is uncertainty about both the actual cost for compliance and the specific derogations for small-scale providers. SMEs and startups should look to guidance from member states as well as local startup associations for how to leverage this support.

Against this backdrop, a structured and data-driven analysis of how German AI startups are distributed across these risk categories becomes indispensable. Only by grounding regulatory discussions in empirical evidence can policymakers and industry stakeholders design targeted support measures that balance compliance with innovation.

4 For further detailed information on regulatory sandboxes including their relevance to AI startups, please see (Fetic *et al.*, 2025)

3. Goals, Methods & Data

3.1 Study Goals

While an impact assessment conducted by the European Commission has estimated a 5-15% share of high-risk AI systems, there are no metrics for the German startup ecosystem regarding the distribution of risk classifications.

Accordingly, **a key goal of our study was to understand:**

> ***What is the distribution of German AI startups across EU AI Act risk tiers, and how does it compare to prior policy assumptions?***

Based on additional data available to us from the German AI Startup landscape Survey, **we also attempted to understand:**

- **Industry Sectors:** Which industries and enterprise functions concentrate high-risk and transparency obligations?
- **Länder:** How do risk-tier distributions and sectoral compositions vary across German federal states?
- **Funding:** What does the funding pattern look like for startups across the various risk classes?

3.2 Methods and Data

Accordingly, we examine the regulatory landscape of German AI startups under the EU AI Act, aiming to provide a comprehensive understanding of their distribution across risk categories and the implications for compliance and oversight. We analyze how risk classifications vary depending on technology types, AI capabilities, and enterprise functions. Geographic and financial dimensions are also considered, highlighting patterns in where high-risk or transparency-obligated startups are concentrated and how investment dynamics intersect with regulatory requirements.

To do this, we compiled a dataset of 628 German AI startups and 2,451 AI use cases based on publicly available descriptions (websites and product materials). Each use case was classified under the EU AI Act risk framework using an LLM-assisted classification pipeline, followed by expert legal review and manual checks. For each startup, we report the highest-risk classification across its portfolio to reflect worst-case compliance exposure. Additional methodological documentation, including dataset construction, model/prompt setup, classification decision rules, and validation results, is provided in Annexes I and II.

4. Results & Findings

This section presents the empirical findings of our study: A structured overview of how German AI startups' AI use cases are distributed across the EU AI Act's risk categories. These form the foundation for assessing sector-specific implications and identifying areas where targeted support and capacity building will be most needed.

We structure each finding in the following manner:

- Key **results**
- Policy **implications** of our findings for the EU Commission, the German federal government & states, and for startups and the broader industry
- **Recommendations** for these stakeholders
- **Limitations** inherent to our study and specific to the finding.

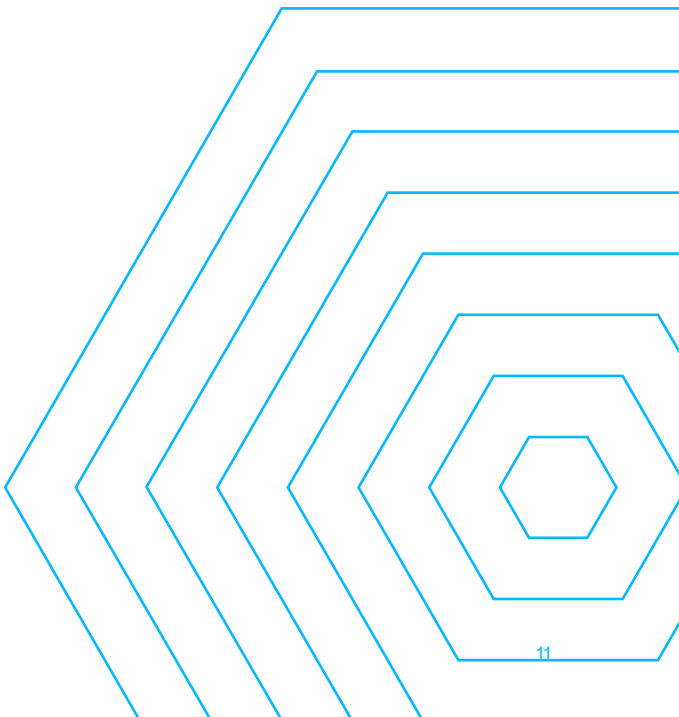

Finding 1: The share of high-risk AI startups significantly exceeds the European Commission's original assumptions

Figure 1 shows the frequency of each risk classification in the results. Based on the highest-risk AI use case per startup, 25.4% (142 out of 559) of the analyzed startups are classified as high-risk. A further 26.3% (147 startups) are subject to transparency obligations, while 48.3% (270 startups) fall into the low-risk category. Notably, none of the analysed startups were found to develop prohibited AI systems.

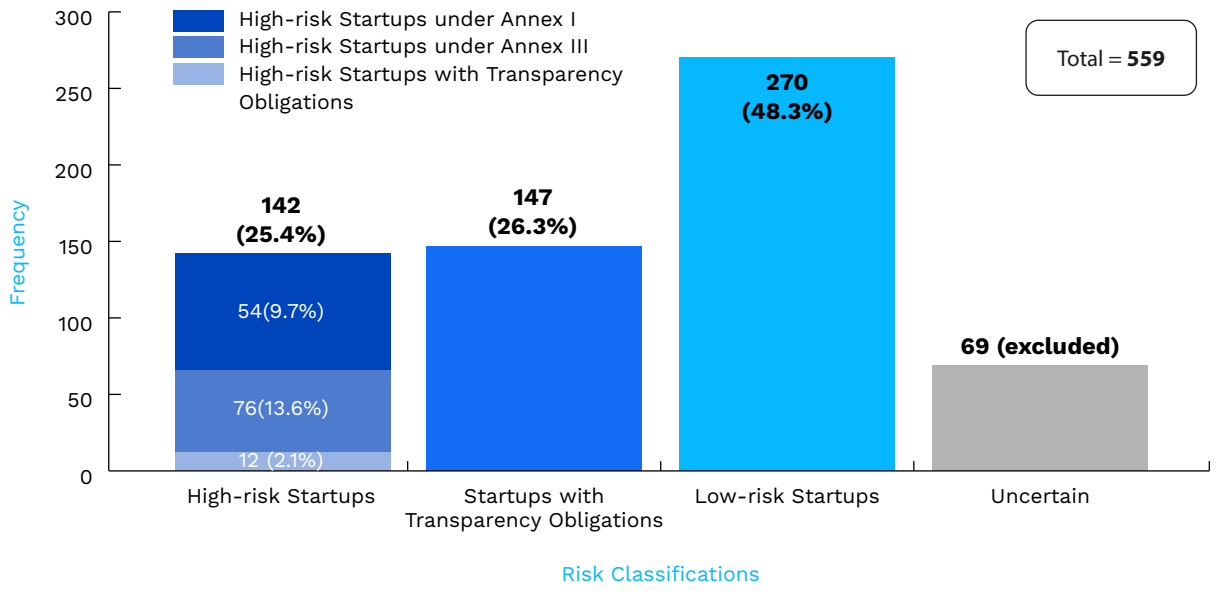

Figure 1: Distribution of Risk Classification of AI Startups

Implications

Stakeholder	Implications
EU Commission and the German Federal Government	• The EU Commission's 2021 impact assessment, which estimated that 5%–15% of use cases would be high-risk, does not apply to German AI startups.
States / Regions	• Market surveillance authorities will likely face a higher volume of high-risk systems than initially planned. • Given the extensive requirements under the draft Code of Practice (CoP) for transparent AI systems (e.g., training, tools, documentation), market surveillance authorities will also likely have to monitor a larger-than-expected volume of AI systems with transparency obligations (see Finding 3 for additional details).

Recommendations

Stakeholder	Recommendations
EU / AI Office	• Revisit the 2021 impact assessment, including the compliance costs associated with the AI Act, and encourage member states to conduct similar studies in their jurisdictions. • If necessary, adjust supportive measures, including funding mechanisms, to account for “the real/observed” regulatory burden.
German Federal Government / States / Regions	• Competent authorities and market surveillance authorities need to be sufficiently equipped to cope with the volume of regulated AI systems (e. g. Bundesnetzagentur) • Ensure that startups with low-risk use cases are not marginalised. They still face significant regulatory burdens from other frameworks (eg, GDPR, liability, and IP), and should be considered for measures for fostering AI innovation.
Industry / Startups	• Integrate AI Act risk classification into early-stage product planning to avoid "costly surprises" and the potential need for redesigns later. • Investors should integrate regulatory compliance costs into valuation models and funding rounds in order to ensure that German startups are competitive and do not divert resources from product to compliance.

Limitations

- **Scope of the analysis:** Our findings present evidence for German startups and may not generalise to other countries or enterprise sizes.
- **Information basis:** We scraped the use-case descriptions from startup websites and summarised them using LLMs. These descriptions may be incomplete, simplified for marketing purposes, or insufficiently precise to fully capture the deployment context or regulatory implications.
- **Methodology:** This classification combines automated LLM analysis and human review. While this enables scalability, both have a margin of error.
- **Granularity of review:** The analysis does not disaggregate sub-categories within Annex I or Annex III of the EU AI Act, nor does it differentiate between individual transparency obligations under Article 50.
- **Interpretative limitation:** Given the absence of guidance, we did not consider whether any use cases might benefit from derogations in Article 6(3) of the AI Act.

Finding 2: Context and missing details limit fully automated risk classification and complicate human review

A total of 69 startups could not be conclusively classified due to incomplete information about their use cases (even after manual evaluation of all such use cases). These startups are excluded from further analysis. **Table 1** provides examples of use cases that we were unable to classify and the source of our uncertainty.

Examples of use cases where uncertainty about interpreting the law or a lack of information about the AI system prevented classification

Example 1	Example 2	Example 3
Battery Analytics	**Data Aggregation**	**Motion Capture**
Intended Purpose: A predictive battery analytics platform designed to enhance the safety, performance, and longevity of lithium-ion battery systems.	**Intended Purpose:** An AI agent that continuously scans and analyzes global news sources to deliver real-time intelligence on current events and developments.	**Intended Purpose:** The system provides highly accurate, real-time tracking systems and motion analysis of humans based on simple video recordings and is deployed across various sectors, including sports organizations, and medical clinics.
Uncertainty: The risk classification depends heavily on the context of use. If such a system performs a safety function, it could potentially be high-risk due to use in critical infrastructure or due to use in products regulated under Annex I legislations.	**Uncertainty:** It is unclear from the use case description if the system generates synthetic content (which would create transparency obligations) or merely identifies relevant news content.	**Uncertainty:** It is unclear if the system is also capable of identifying target individuals.

Table 1: Examples of Uncertain Classifications

Implications

Stakeholder	Implications
EU / AI Office	• The AI Act's assumption that every use case will have one "intended purpose" might not always hold. Risk classification can shift based on deployment context and sector. In fact, our study shows that startups often operate across sectors (cf. Finding 4 for details). • Ambiguous risk classification can result from various causes, including multi-use AI systems, incomplete or ambiguous use-case descriptions, or inaccurate interpretations of the AI Act.
German Federal Government/ States / Regions	• In the absence of clear guidance, decentralised oversight might lead to conflicting interpretations of the AI Act and of the risk classification results between oversight bodies.
Industry	• The same AI system can fall into different risk classes based on the application sector or enterprise function (e. g., document analysis in human resources might be high-risk vs. low-risk in marketing).

Recommendations

Stakeholder	Recommendations
EU / AI Office	• Provide clarifying guidance on how "intended purpose" should be interpreted in general, including for multi-purpose or modular AI systems or multi-agent systems.
German Federal Government / States / Regions	• Adopt a guidance-first approach to market surveillance that rewards good faith efforts and acknowledges the technical challenges with risk classification. • Establish coordination mechanisms to ensure consistent interpretation of risk classification across supervision authorities. • Support local initiatives that help clarify the risk class and can remove legal uncertainty for industry and the public sector.
Industry / Startups	• Deployers and providers should be mindful about the specific purpose and context of use for each AI system, as this may have consequences in the form of additional/other obligations under the AI Act[5]. • Product teams in startups should be conscious about changes to the product or use case, as iterations might change the risk classification.

Limitations

- **Information basis:** We scraped the use-case descriptions for startup websites and summarized them using LLMs. These descriptions may be incomplete, simplified for marketing purposes, or insufficiently precise to fully capture deployment context or regulatory implications.
- **Interpretative uncertainty:** Legal interpretation of several AI Act concepts remains unsettled and may evolve through guidance or case law. Until then, human reviewers may have misunderstood the law.

5 See Article 25 of the AI Act

Info Box

The Crucial Role of Human Review in Risk Classification

Accurately classifying AI systems under the EU AI Act is not a purely mechanical exercise. While guidance materials and interpretative resources are increasingly available, our analysis shows that reliable classification requires the combined expertise of both legal professionals and technological specialists. This is because the risk level of an AI system often depends on fine-grained details of its actual deployment context and the technical details of the product, which cannot be captured by general-purpose guidelines alone.

In practice, we repeatedly encountered borderline cases where identical technical functionalities could fall into different categories depending on the application sector or end-use. For example, predictive maintenance solutions for manufacturing machinery may be considered lower-risk, whereas the same approach applied to critical infrastructure, such as wind turbines, may trigger a high-risk classification. Similarly, HR-related AI tools frequently raise uncertainty. For example, AI solutions for shift planning or workers management might be low risk if it is based on employee neutral parameters, such as predicted customer flow, but might easily become high risk if they are based on employee characteristics.

Such challenges are compounded by the AI Act's reliance on references to other EU legal frameworks. Annex I, for instance, incorporates obligations from multiple sectoral regulations, creating interpretive grey zones. Robotics illustrates this complexity: not all AI-powered or enabled machines must undergo a third party assessment under the current machinery directive; some might have to under the forthcoming machinery regulation; while others yet might be governed by other sectoral laws, such as those relating to vehicles or agricultural devices.

These examples highlight the inherent limits of desk research based solely on publicly available information. In our study, a number of use cases could not be conclusively classified due to missing deployment details or unresolved legal interpretations. This underscores the crucial role of humans in the risk classification process and the limitations of automated approaches or over-simplified checklists. Human judgment—anchored in both legal interpretation and technical expertise—remains indispensable for ensuring accurate and context-sensitive application of the EU AI Act for the time being.

Finding 3: The diffusion of GPAI models is making transparency obligations the dominant regulatory baseline for AI startups

Our German Startup Landscape Study 2025 showed that nearly one in three German startups are building generative AI solutions, with the number of such startups growing 130% over 2024. Unsurprisingly, 59% of generative AI startups in our study—which is based on the 2024 German Startup Landscape Study—would likely be subject to transparency obligations. Moreover, 12% of these startups were also classified as high-risk. **Figure 2** shows the distribution of a total of 111 Generative AI startups across each risk category.

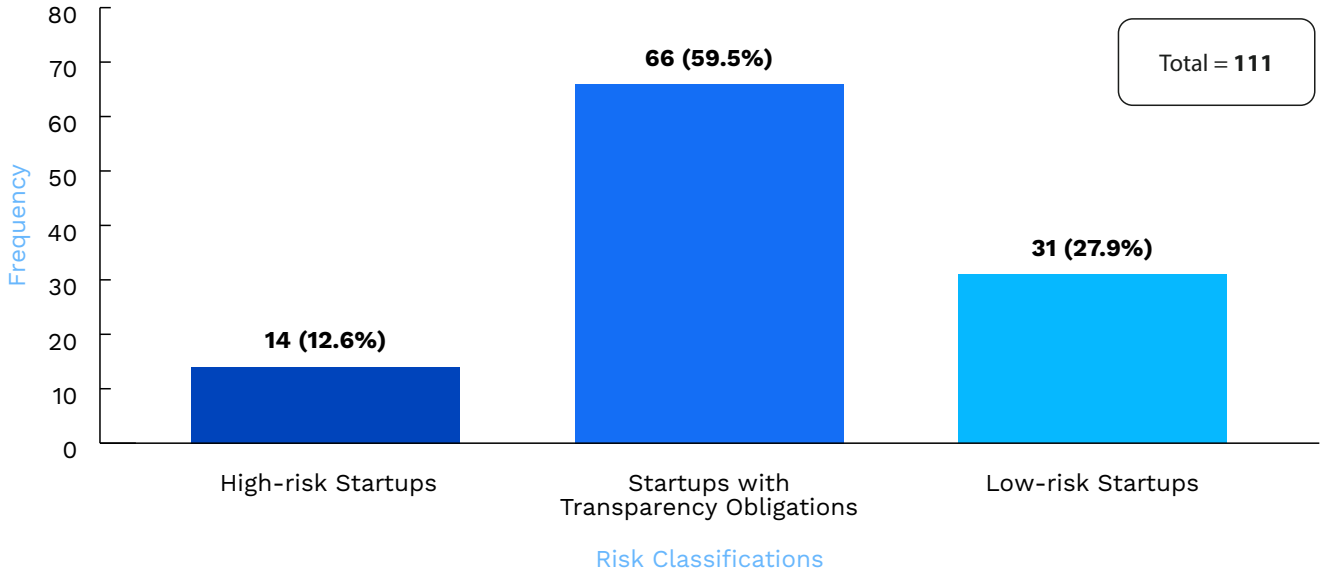

Figure 2: Risk Classification by Generative AI startups

Implications

Stakeholder	Implications
EU / AI Office	• Transparency obligations are likely to affect a larger (and potentially growing) share of German startups as LLM use increases over time. • The ability of these startups to comply with the AI Act, particularly the high-risk requirements, will be significantly implicated by the extent to which GPAI providers adhere to the information and resource sharing obligations under the code of practice for general purpose AI models.
Industry / Startups	• The performance and regulatory compliance of German AI startups is increasingly tied to non-EU general purpose AI model providers. AI startups need to weigh the potential dependencies they accept when choosing a GPAI model as the basis for their AI system or business model.

Recommendations

Stakeholder	Recommendations
EU / AI Office	• Consider adopting the Omnibus proposal to delay the transparency obligations, given that the Code of Practice for transparent AI systems will only be finalised in summer 2026. • Monitor the functioning of the Code of Practice for general-purpose AI models and ensure that downstream providers receive sufficient information to ensure compliance. • Support the development of EU based general purpose AI model providers to ensure end-to-end compliance along the AI Value Chain.
Industry / Startups	• Ensure that technical details required for downstream compliance are provided by the general-purpose AI model provider as part of model selection criteria, particularly when the use case is high-risk.

Limitations

- **Definition of a "generative AI" startup:** We rely on the classification process used in our German Landscape Study 2024 to determine when a startup qualifies as a generative AI startup and did not develop our own method.

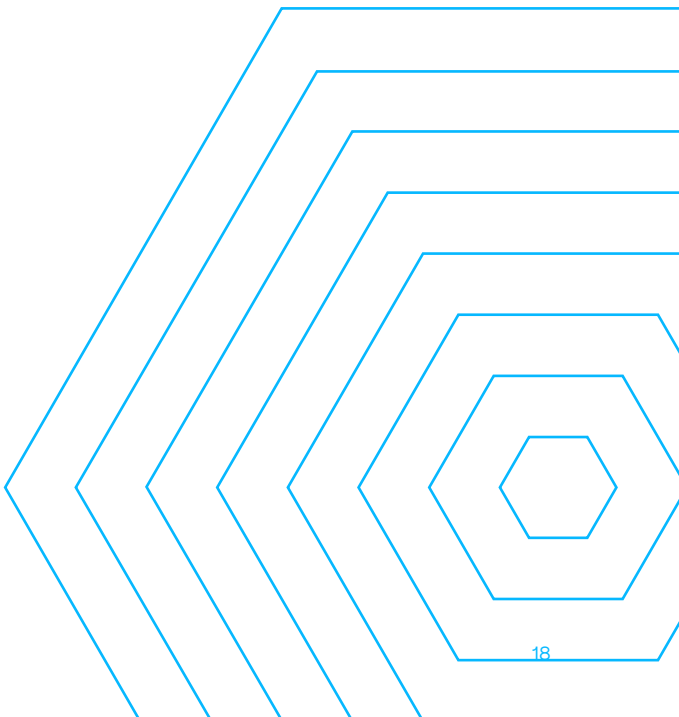

Finding 4: The burden of compliance is concentrated on certain industry sectors and enterprise functions

The compliance burden for German startups is not evenly distributed. High-risk and transparency obligations are often concentrated in specific industrial sectors and enterprise functions.

Figure 3 shows the frequency of all risk categories across each industry. The image displays data for 369 startups for which we were able to manually assign an industry type.

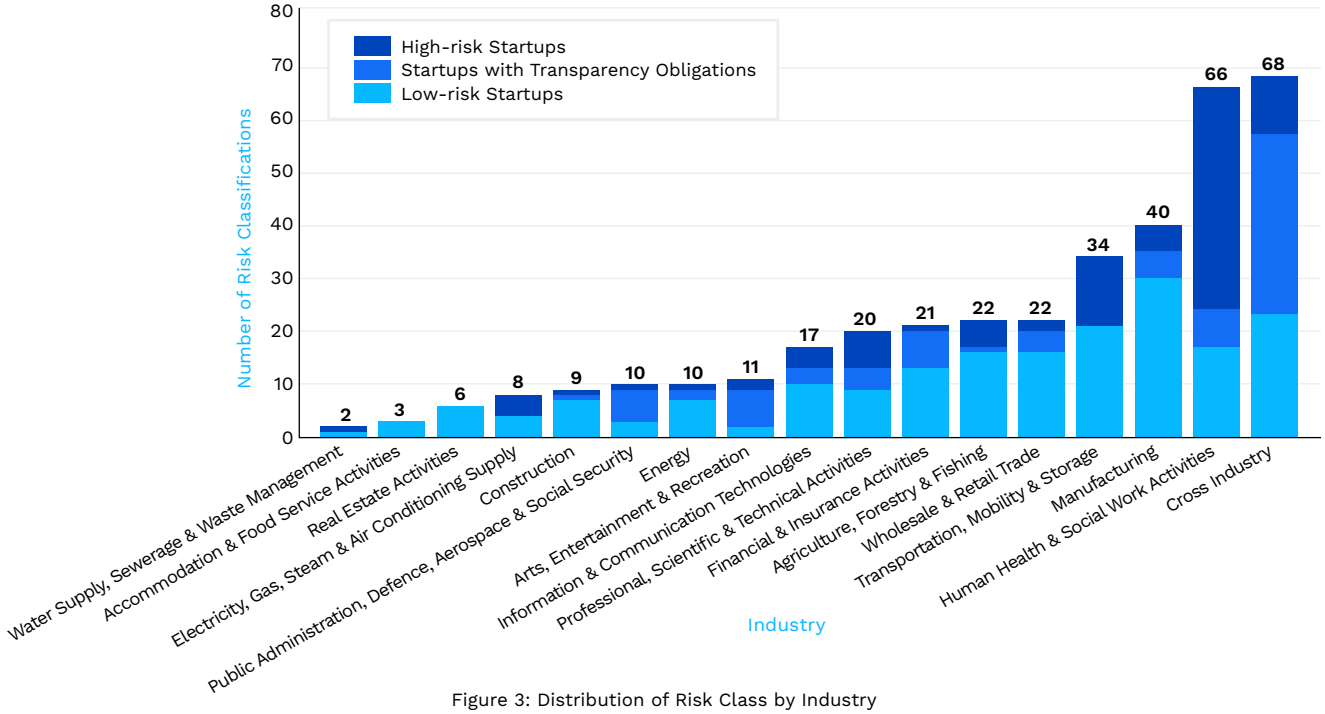

Figure 3: Distribution of Risk Class by Industry

We see a concentration of high-risk startups in industries such as human health (e.g., AI in medical devices), transportation (e.g., autonomous driving) and in “cross-industry” use cases (typically startups whose product is industry-neutral, such as human resources). Similarly, we notice that a significant share of startups that operate across industries likely face transparency obligations.

Figure 4 shows the distribution of risk classifications by enterprise function. The graph shows data for 282 use cases for which we were able to manually assign an enterprise function.

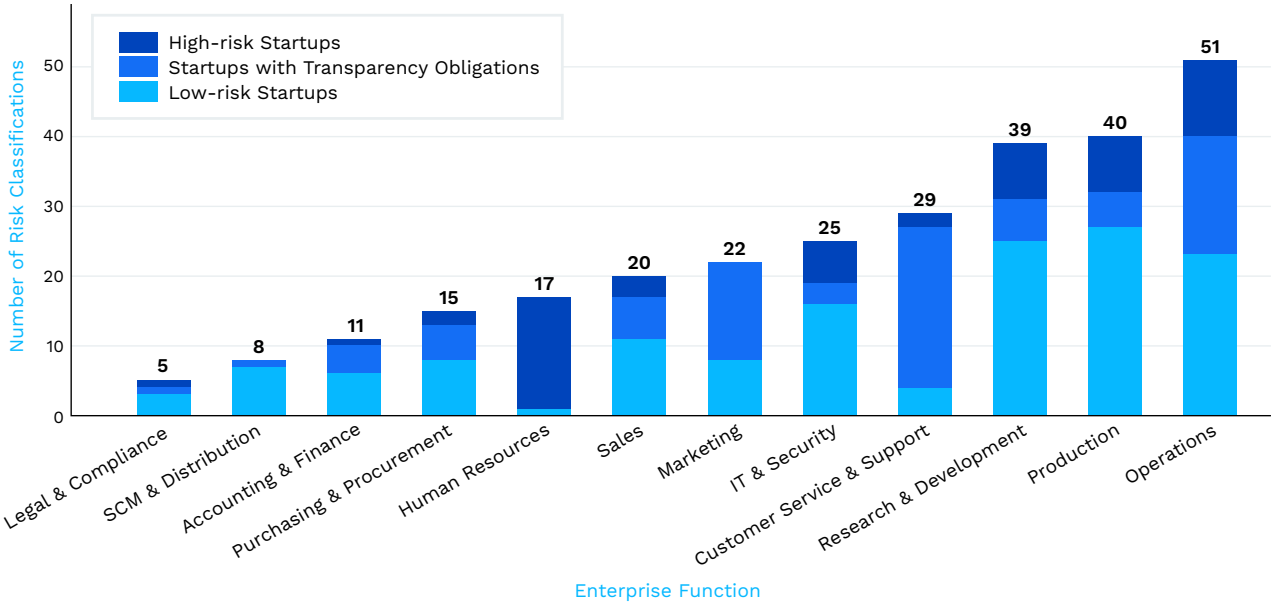

Figure 4: Distribution of Risk Class by Enterprise Function

The highest concentration of high-risk startups is found in enterprise functions like human resources, reflecting a growing trend of using AI systems to bolster hiring capabilities. Similarly, we see transparency obligations concentrating in functions like marketing and customer support where LLM-based AI systems are increasingly used by companies. Enterprise functions with high shares of low-risk AI systems include Production, R&D, and Operations, thus indicating areas with no additional requirements posed by the AI Act.

Implications

Stakeholder	Implications
EU / AI Office	• Uniform and "horizontal" (i.e., industry-agnostic) guidance on compliance might not reflect the reality that regulatory burdens are concentrated in specific industry sectors and enterprise functions.
States / Regions	• Some sectors and domains will require disproportionately more regulatory support and oversight from market surveillance and notified bodies.
Industry / Startups	• Some industry sectors and enterprise functions are structurally more likely to be high-risk. For example, use cases in HR and operations demand more oversight.

Recommendations

Stakeholder	Recommendations
EU / AI Office	• Encourage the development of industry specific guidance on the intersection of the AI Act and other regulations, such as medical devices.
States / Regions	• Compliance support will have to be tailored according to industry and enterprise function. For instance, many companies might benefit from guidance on "AI Act in Human Resources" • Conduct further research about AI-first SMEs in the federal states to have a more complete picture of sectoral risk profiles.
Industry / Startups	• Prioritize upskilling and retraining of employees developing and using AI systems. • State-level Industry associations and SME-hubs should collaborate to reduce the costs of legal interpretation and technical compliance.

Limitations

- **Scope of the analysis:** Our findings present evidence for German startups and may not generalise to other jurisdictions or enterprise sizes.

Finding 5: German States exhibit similar profiles on risk-class, but different distribution across industrial sectors

Across German states with significant startup activity, the distribution of AI Act risk classifications is broadly similar despite large differences in the absolute number of startups. However, the composition of startup activity across industrial sectors differs markedly between states, leading to different patterns of regulatory exposure.

Figure 5 shows the distribution of risk classification by federal states and **Figure 6** narrows in on the distribution for the top five states by total number of use cases.

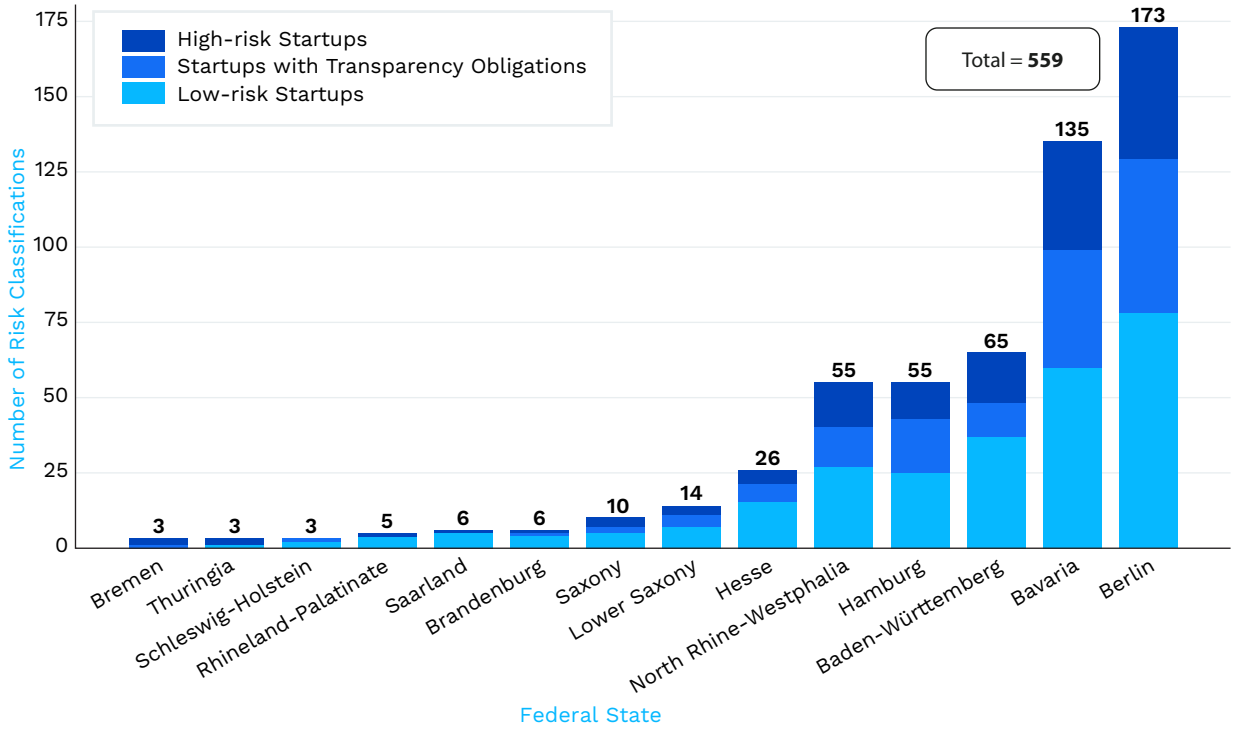

Figure 5: Distribution of Risk Classification by Federal States

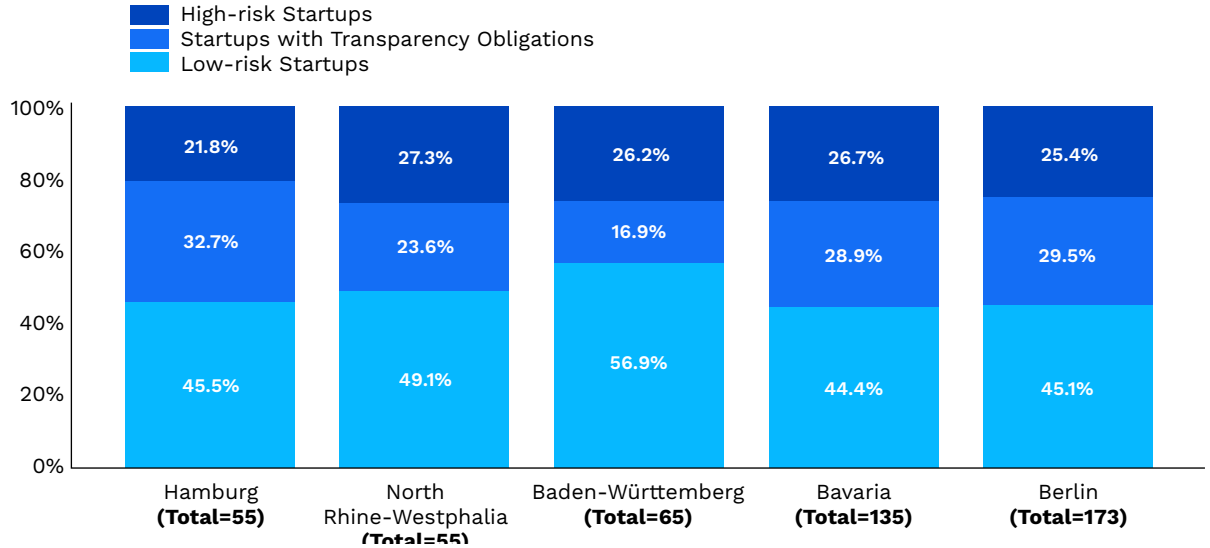

Figure 6: Distribution of Risk Classification by Federal States with Highest Startup Activity

Among the five federal states with the highest number of AI startups, the distribution of startups that are low-risk, have transparency obligations, and are high-risk is broadly similar. For the remaining federal states, total startup activity is too low to draw any meaningful comparisons.

In comparison, sectoral distributions look very different. **Figure 7** shows a heatmap of the 351[6] AI startups across states with at least eight use cases and industries for which manual encodings are available.

The number in each cell denotes the number of startups in the industry labeled on the x-axis and located in the state labeled on the y-axis. The color scale in the heatmap runs from light to dark blue, where light blue represents a low number of startups and dark blue indicates a high number. Cells without any startups remain blank.

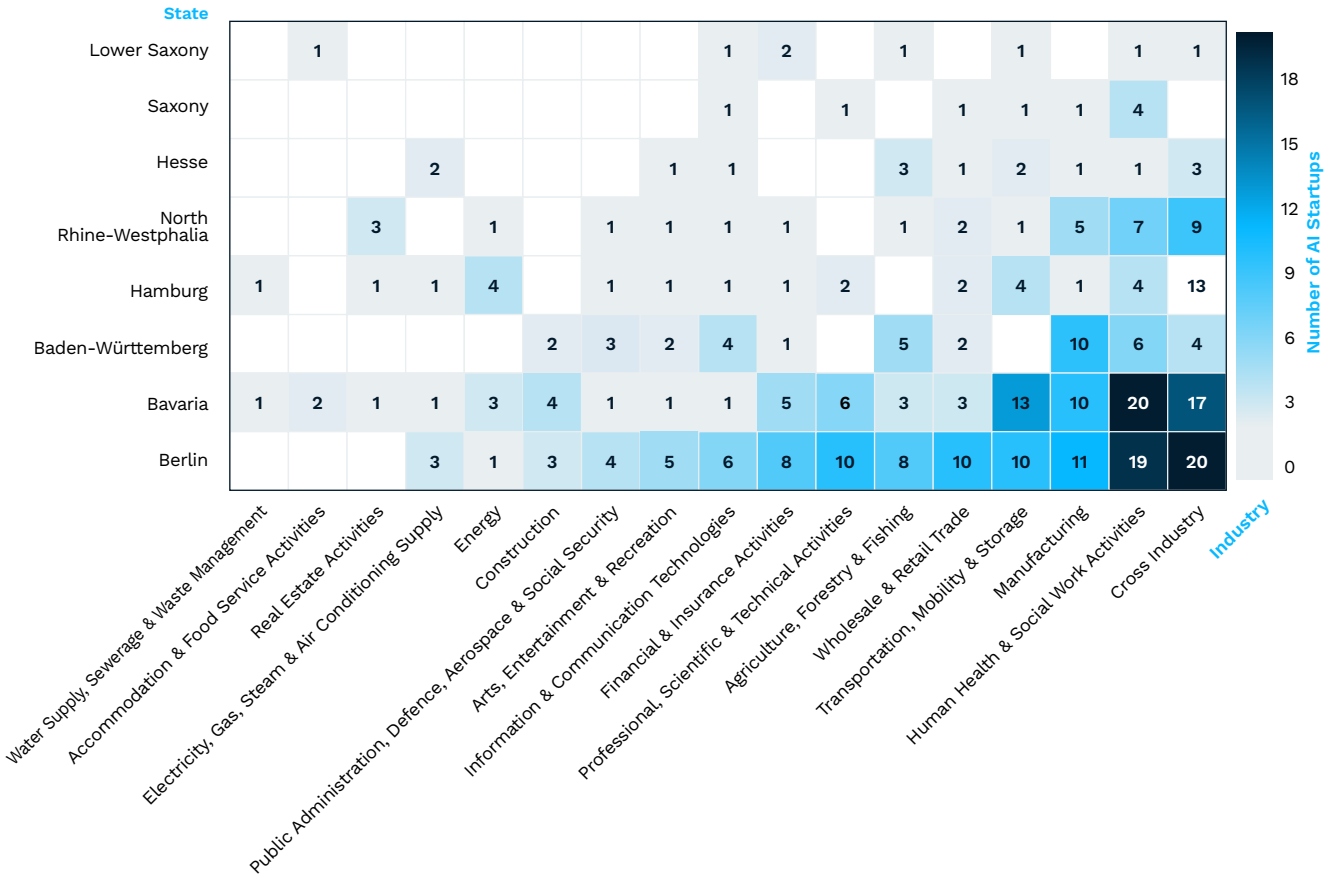

Figure 7: Heatmap of AI Startup Activity by Federal State and Industry

The states towards the lower end of the y-axis (e.g., Bavaria and Berlin) and industries at the right end of the x-axis (e.g., Human Health and Cross Industry) have a higher number of startups. Bavaria is the only state represented across all industries, while Berlin is home to the highest number of total startups. Both states show high concentrations in cross-industry, human health, manufacturing, and transportation industries. Other states show significantly narrower industry profiles, often concentrated in manufacturing, energy, real estate, or other industry sectors.

6 Eight federal states had fewer than eight startups. We excluded them from this analysis. For the remaining states, we only included use cases for which industry encodings were available (See also Fig 3).

Implications

Stakeholder	Implications
States / Regions	• Regulatory needs differ substantially across states, requiring differentiated supervisory, advisory, and support capacities.
Industry / Startups	• Many startups operate in the “cross-industry” sector, i.e., they build industry-agnostic products and services.

Recommendations

Stakeholder	Recommendations
States / Regions	• Establish cross-state cooperation mechanisms focused on industry-specific AI use cases, enabling the sharing of expertise, supervisory resources, and best practices. For instance, manufacturing in Berlin, Bavaria and Baden-Württemberg. • Commission further empirical analysis of state-level industry composition to better align supervisory capacity and support measures with regional risk profiles. • Anticipate that regulatory demand might not be the same across Germany and ensure that oversight and market surveillance bodies are appropriately resourced.
Industry / Startups	• Strengthen structured support for AI startups through industry associations and private capital actors, particularly in high-risk or cross-sectoral domains, to improve access to compliance expertise and shared governance resources. This support could be offered through institutions like, for example, the VDMA, ZVEI, BDI, or BVMED..

Limitations

- **Small sample sizes:** For many federal states, the total number of startups in our study is small, limiting meaningful comparisons.
- **Partial ecosystem coverage:** The analysis covers AI startups only. As a result, the findings do not represent the full regional AI landscape, e.g. the prevalence of AI in SMEs and corporations.

Finding 6: Compliance cost is not (yet) reflected in their funding

Despite a substantial share of startups classified as high-risk or as having transparency obligations, funding patterns do not yet reflect possible compliance costs associated with the AI Act.

Figure 8 shows the average and median funding[7] per risk classification. Average funding is highest for transparency-obligated startups and lowest for low-risk startups. The much lower medians compared to the means across all groups show strong right-skew. In other words, a few outlier startups receive a significantly larger share of funding than the remaining majority.

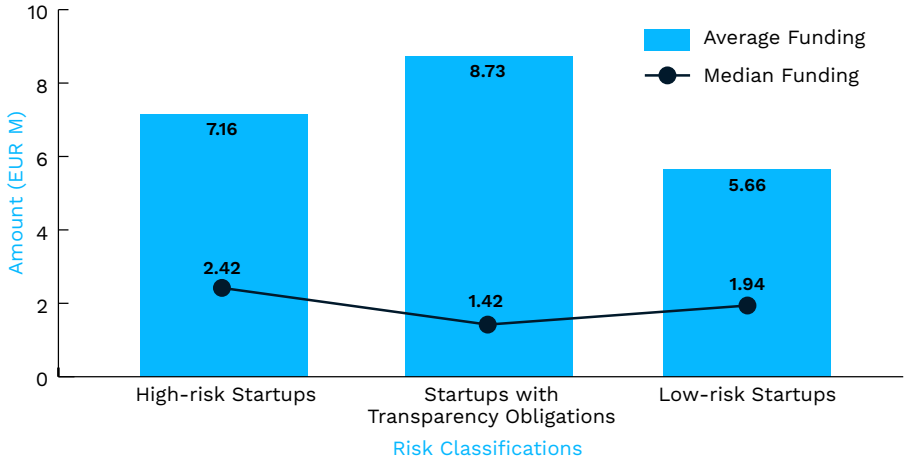

Figure 8: Average and Median Funding of Startups by Risk Classification

Figure 9 shows the total funding received by all startups in a given federal state. The top five federal states by total funding are the same states with the highest number of AI startups as shown in Figure 5, and in the same order, i.e., Hamburg, North Rhine-Westphalia, Baden-Württemberg, Bavaria, and Berlin. Finally, **Figure 10** shows the distribution of total funding across risk categories of these five states with the most startup activity.

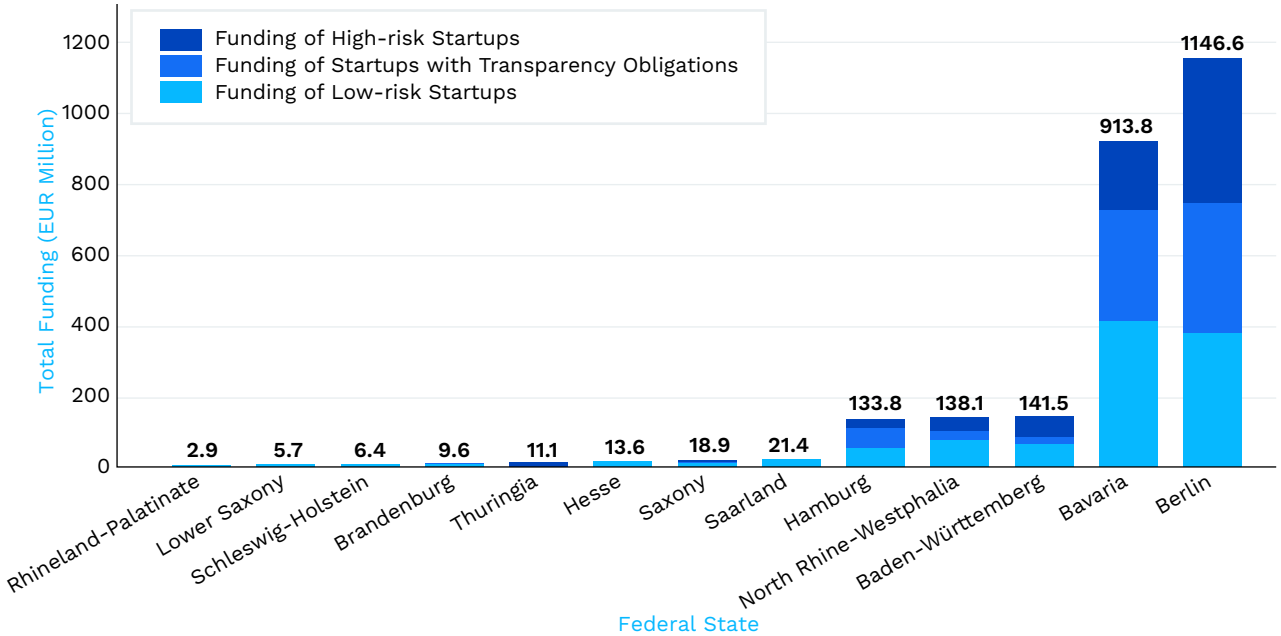

Figure 9: Total Funding by Federal State and Risk Classification

7 The data was sourced from dealroom.co and complemented with internal data sources in July 2024. It represents total capital raised, primarily from private market investors.

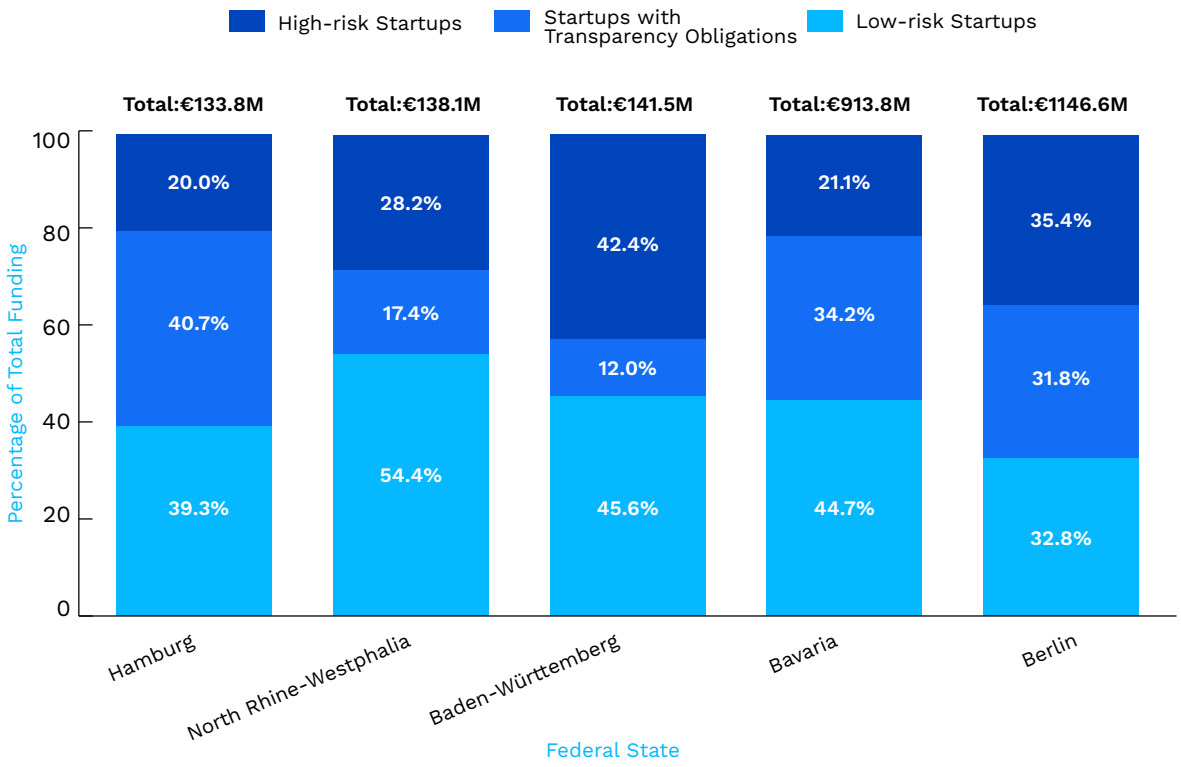

Figure 10: Funding Distribution by Risk Classification for States with the Most Startup Activity

Given the right skewed nature of funding (i.e, a small percentage of startups receive the largest share of funding) and the comparatively low number of AI Startups from federal states apart from Berlin and Bavaria, we believe that this data should be interpreted with caution and not as indicating structural patterns across investments (see the limitations below).

Implications

Stakeholder	Implications
States / Regions	• The combination of the modest median per startup funding across all risk classes (see Figure 8) and the expected compliance cost of approximately 200k-300k Euro for reaching compliance of one High-Risk AI System (Renda *et al.*, 2021, pp. 152–153), allows an approximation of the relative regulatory burden for startups.
Industry / Startups	• Highly skewed funding in all classes implies that well funded startups might cope better with compliance as opposed to the many startups with modest funding • The tolerance for compliance cost by VCs likely differs across sectors (e.g. biotech vs. EducationTech), which might influence the chance to attract funding for AI Startups and thus the survival of such startups.

Recommendations

Stakeholder	Recommendations
States / Regions	• To lower the dependency of a startup's survival on coping with compliance costs, state governments (with high regulatory burden) should offer ample support. • Commission further research to study if risk classification affects funding requirements and by exactly how much.
Industry / Startups	• Further studies are needed to understand the willingness of VCs to fund high-risk AI startups.

Limitations

- **Causation**: We have not investigated all of the variables that go into startup funding. Regulatory compliance is only one factor and a high-risk classification by itself does not necessarily warrant greater funding.

5. Conclusion

This study delivers the first comprehensive and empirically grounded risk classification of the German AI startup ecosystem after the passing of the EU AI Act. Our findings reveal a greater than expected regulatory impact for German AI startups, with roughly 25% of startups potentially being high-risk and another 26% potentially subject to transparency obligations. To prevent the AI Act from becoming a barrier to innovation, we call on EU, German, and industry stakeholders to, among other things, improve guidance, enhance regulatory capacity, and simplify compliance. We hope that these insights spur further research into compliance costs and regulatory implications for startups and inspire actors across the AI Ecosystem to take invested action for accelerating trustworthy AI Innovation across Europe. Continued data-driven analysis of the AI Act, including its impact on AI innovation, will continue to be an essential pillar for evidence-based decision-making toward reaching Europe's AI ambitions.

Annex

Annex 1: Dataset

This study set out to consider all AI startups which are part of the German AI Startup Landscape 2024 (n=687)[8]. 69 startups from the landscape were identified to be no longer eligible due to inactive websites, liquidation, acquisition, or moving their headquarters out of Germany, and were hence removed from this study. There are two ways to join the German AI Startup Landscape. Startups can either apply to be part of the landscape via our application form, or they can be nominated by one of our partners, those mainly being venture capital firms and tech companies.

To be accepted into the AI Startup Landscape, each startup needs to fulfil the following criteria, which are evaluated by our AI experts:

- Be a registered company.
- Have been founded less than 10 years ago or pivoted their core business model to AI less than 10 years ago (including 2014).
- The startup's headquarter must be situated in Germany.
- They should have a minimum of two Full Time Equivalents (FTEs)
- They should have a minimum of one FTE with AI competence
- They should have AI at their core or exhibit a significant usage of AI.
- The startup respectively its business model has a high face validity (e.g., professional website, convincing business model, etc.).

For detailed information about the underlying data (e.g., with regards to industries, underlying technologies, and much more), we refer to the detailed analysis of the AI Startup Landscape[9]. Additionally, detailed funding data on the AI Startups was sourced from Dealroom and enriched with further internal data.

8 A handful of startups from the AI Startup Landscape 2024 were excluded for the purpose of this study for various reasons such as acquisition, liquidation, and others.

9 AI Startup Landscape 2024: https://www.appliedai-institute.de/en/hub/2024-ai-german-startup-landscape

Annex 2: AI Risk Classification Tool

During the data collection, all respective use cases the startups offered were classified along the risk categories of the EU AI Act: unacceptable risk, high risk, transparency risk, and low risk (see Chapter 2.1 for more info). Accordingly, we proceeded in three steps. First, the automated extraction of use cases from the startups in the German AI Landscape. Second, an ensemble-based classification using leading LLMs. Finally, a manual human review of the risk classifications.

Specifically, we developed an automated and prompt based AI risk classification tool which in the first step collected use cases of all 628 AI startups eligible for evaluation using OpenAI's Web Search API. In a second step, using ensemble-based classification with LLMs each extracted use case was analysed and classified individually and our tool provided detailed reasoning for each decision. The classification prompt, which can be found in the GitHub repository, was drafted by our AI experts and lawyers. It compiles the AI Act in a machine processable manner, giving clear instructions for the automated risk assessment. Creating it was part of an iterative process together with the technical development (more information below). In a third step, our expert team conducted a manual human review of every single risk classification (cf. Insight Box "The Role of Human Review in Risk Classification").

Obviously, startups often offer more than one solution, i.e. they have several use cases for their solution. For the purpose of this study, we report the highest risk classification of each startup. For example, if a startup offered both a low-risk AI-powered document summarization tool and a high-risk predictive healthcare analytics system, we chose the startup based on the higher-risk category. This approach ensures that we account for the most critical regulatory compliance considerations associated with each startup's offering portfolio. To guarantee the quality of the automated categorization, our experts cross-examined all entries and validated the output manually.

Workflow

An application was developed in Python to extract AI use cases of the startups in the AI Startup Landscape and to classify each identified use case according to the rules defined in the classification prompt. The code is publicly available on GitHub. The architectural overview of the tool is outlined in the diagram (**Figure 11**) below:

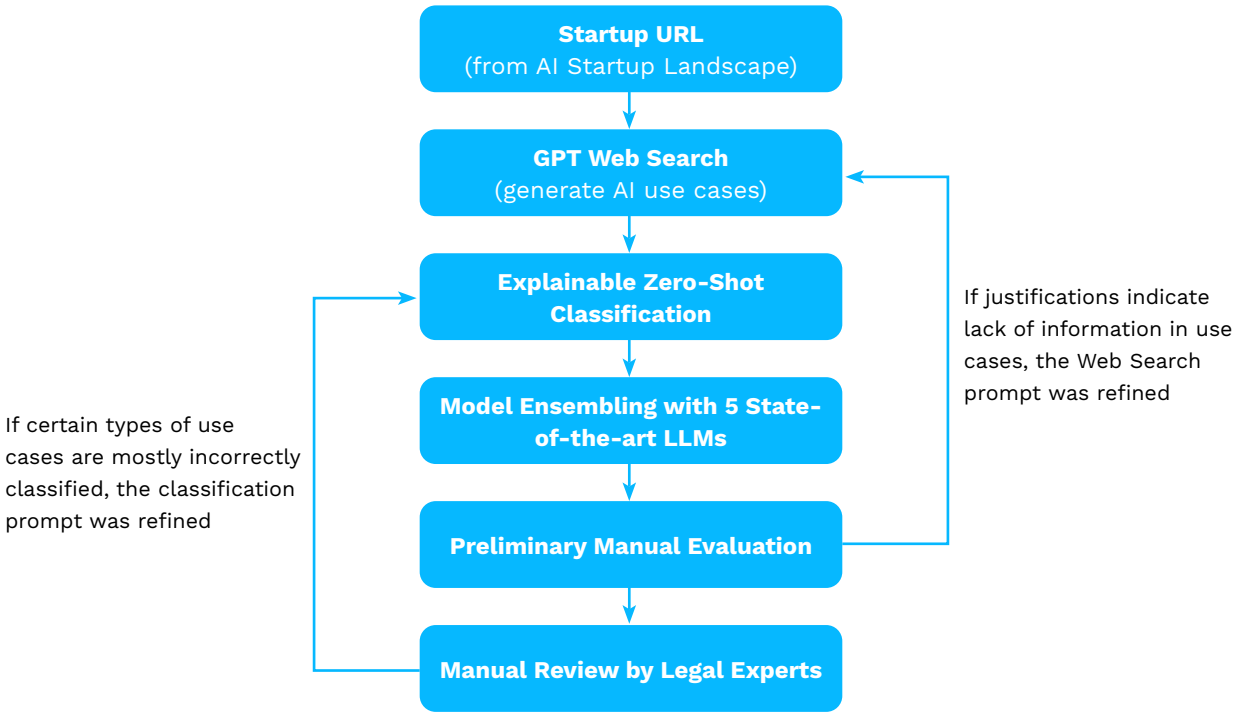

Figure 11: Application Architecture

The first step was to extract all the AI use cases of startups. At the time of conducting this study, OpenAI released a "Web Search" API for a limited number of models, specifically "gpt-4o-search-preview" and "gpt-4o-mini-search-preview. While the former model is larger, the so-called "search context size" parameter in the API controls how much content is retrieved from the web. The API provides three values for this parameter: high, medium, and low, with decreasing order of search context size. For our purpose, we opted for the gpt-4o-search-preview model and chose the highest search context size so that the maximum number of web pages could be scraped in order to get the best possible use cases. While a higher search context size is expensive, the trade-off for quality was justified.

All AI use cases for each startup were stored in a .csv file. The Web Search api was instructed to collect the most relevant information for our zero-shot classification task. This included the following details in the generated use cases:

1. Intended purpose
2. Deployment sector
3. Level of autonomy
4. Impact on individuals
5. Types of data used
6. User type
7. Adaptivity in deployment
8. Safety-critical nature of system

Model Ensembling with majority voting

After all the AI use cases for each startup were obtained, the next step was to classify each use case into the different categories of the EU AI Act. The approach we have used is called "Explainable zero-shot classification". The term "explainable" refers to the fact that the models are prompted to generate a justification along with each classification label[10].

The categorical classes for this purpose were defined as follows:

1. Prohibited AI system
2. High-risk AI system under Annex I
3. High-risk AI system under Annex III
4. High-risk AI system with transparency obligations
5. System with transparency obligations
6. Low-risk AI system
7. Uncertain

The classification prompt contained the precise criteria of assigning a risk category from the above list to each generated AI use case. The input token count was well within the maximum context lengths of all the models used.

Ensembling with LLMs

In conventional machine learning tasks, ensembling is a method where multiple models are used simultaneously to predict a label. While there are several types of ensembling methods such as bagging, boosting etc., the ultimate goal of ensembling is to combine the predictions of multiple

10 Regarding terminology: A Class represents the set of all possible values for a given classification, while a Label denotes the specific value predicted by the model.

models into a single outcome. Generally, this can achieved in a number of ways, and the following are a few examples:

1. Taking a majority vote amongst all the predicted labels
2. Taking the average of all model output values. However, this works with numerical data.
3. Taking a weighted vote or weighted average. Certain models that are known to perform better (i.e. have a higher validation accuracy), get more weight.
4. Select the label that any one model predicts (max voting) or only the label that all models agree (min voting)

For our task, we used the first approach of ensembling with majority voting. For this approach, we selected 5 state-of-the-art LLMs as follows:

1. ChatGPT 4o
2. Claude 3.7 Sonnet
3. DeepSeek Reasoner
4. Gemini 2.0 Flash Thinker
5. Mistral Large

Each AI use case, together with the classification prompt was fed simultaneously into each model using the relevant api. Once all responses from the models are received, the voting mechanism starts. In this process, we look for the most common classification label (or risk level) amongst all predicted labels generated by the models. The following diagram (**Figure 12**) outlines the voting mechanism:

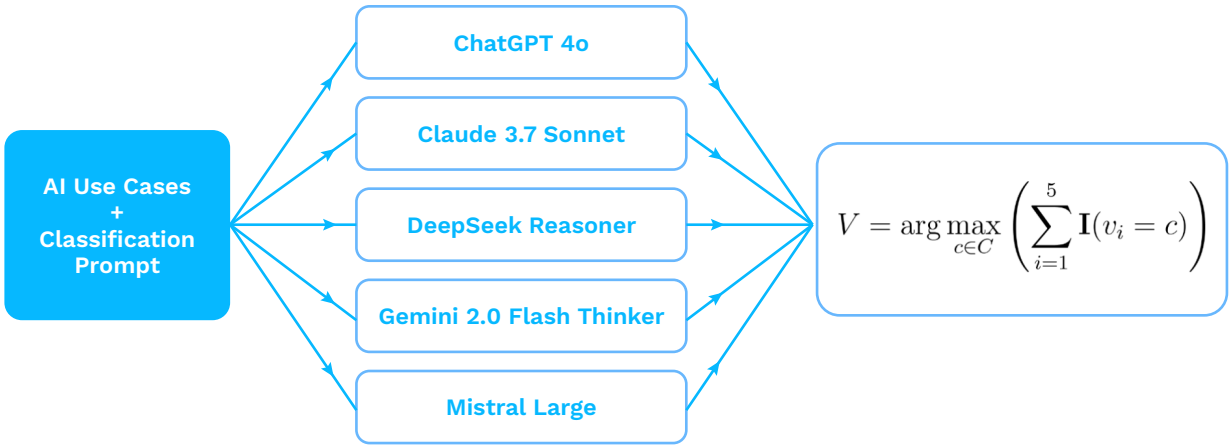

Figure 12: Model Ensembling

Tool Performance

The input starts with a single use case and the classification prompt that is fed into each model simultaneously. Each model generates a response containing the following fields:

1. **Risk Classification:** The classification label for the use case
2. **Reason:** A comprehensive justification for the label that evaluates the use case according to all the instructions in the classification prompt. This reasoning is laid out in 7 separate steps that evaluate the use case against each possible risk class (Prohibited, High-risk AI system under Annex I...Low-risk AI system)
3. **Requires Additional Information:** If the model does not have enough information to make a reasonable classification, the value of this field is set to “Yes”
4. **What Additional Information:** If additional information was required, it is specified in this field

The program waits for each model to generate its set of results. The “Risk Classification” field from each model is then passed to an argmax function as shown above. The function simply counts the number of classification labels generated by each model. This count is referred to as a “vote”. The higher the number of votes a label gets, the higher the chance of that label being assigned to the use case.

For example, consider the following case:

- C as the set of all possible classes
- Set of all votes case by the LLMs (v_i):
  - High-risk AI system under Annex I
  - High-risk AI system under Annex I
  - Low-risk AI system
  - High-risk AI system under Annex I
  - Low-risk AI system
- For each $c \in C$, we compute the following to obtain the majority vote, V:

$$\text{for c = "High-risk...Annex I":} \quad \sum_{i=1}^{5} \mathbb{I}(v_i = \text{"High-risk..Annex I"}) = 3$$

$$\text{for c = "Low-risk AI system":} \quad \sum_{i=1}^{5} \mathbb{I}(v_i = \text{"Low-risk AI system I"}) = 2$$

- For all other classes in C that were not assigned by any of the models, the vote count is set to 0. For instance, “System with transparency obligations” does not appear as any vote v_i. Therefore, the vote for this class will be 0.
- The final majority vote for this example will be:

  $$V = \textit{"High-risk...Annex I"}$$

  This class has the higher number of votes (i.e. 3) and a clear majority. Therefore, it will be assigned as the classification label for the use case.

Once the label is assigned, we have to choose one model generated response from the three responses that won the majority vote in the above example. At this point, all three model responses are correct. Therefore, we simply select the model response that was received first as the final correct answer. The chosen model name along with all other details from the response described earlier are recorded into the output csv.

In the case of a tie in votes, the classification with the lower risk as ordered in this list will be selected. For example for the following votes v_i:

$$\sum_{i=1}^{5} \mathbb{I}(v_i = \text{"High-risk...Annex I"}) = 2$$

$$\sum_{i=1}^{5} \mathbb{I}(v_i = \text{"High-risk...Annex III"}) = 2$$

$$\sum_{i=1}^{5} \mathbb{I}(v_i = \text{"Low-risk AI system"}) = 1$$

There is a tie between “High-risk AI system under Annex I” and “High-risk AI system under Annex III”. The voter mechanism will conservatively pick the lower risk from the tie votes i.e. High-risk AI system under Annex III.

We assessed the performance of the tool via a confusion matrix, visualised in **Figure 13**.

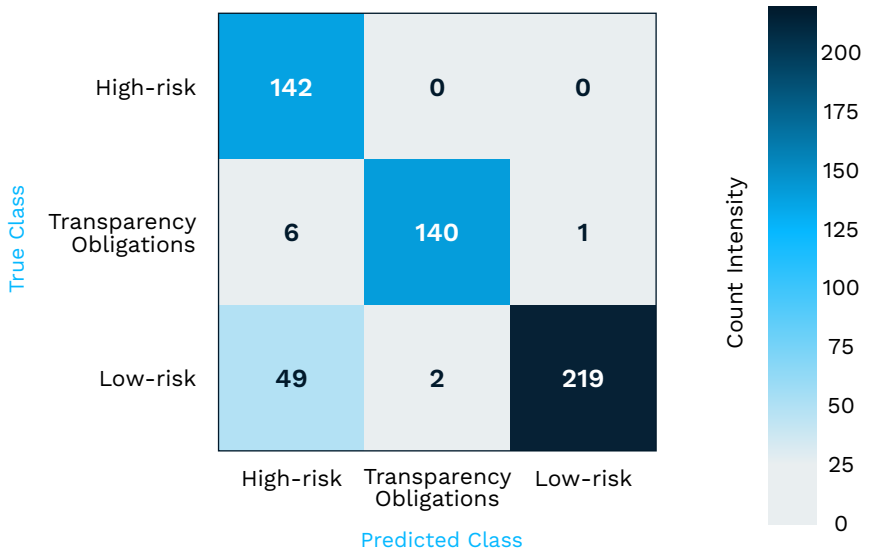

Figure 13: Confusion Matrix

The overall accuracy of the tool can be calculated as follows:

Accuracy = Number of correct predictions / Total number of predictions

= (142 + 140 + 219) / (142 + 0 + 0 + 6 + 140 + 1 + 49 + 2 + 219)

= 501 / 559

= 89.6%

Therefore, 89.6% is the overall accuracy of the tool.

For a six-class legal risk classification task (condensed into three for analysis), the tool shows strong overall accuracy but leans toward caution by occasionally over-classifying low-risk startups as high-risk.

Additional Human Review

The risk labels in this dataset were produced through a hybrid workflow: the ensemble LLMs supplied an initial classification for each of the 600-plus use cases, after which two experienced human reviewers independently checked the output. When both reviewers accepted the machine-assigned label it was recorded as final. Where one or both disagreed, they held an ad-hoc discussion, selected the “most reasonable” class, and—if they felt the models were wrong—over-wrote the entry. We opted against stronger methods because the study’s goal was exploratory mapping rather than regulatory certification.

References

European Commission. (2021). *Commission staff working document: Impact assessment accompanying the proposal for a regulation of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) and amending certain Union legislative acts (SWD(2021) 85 final).* https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=SWD:2021:85:FIN

Hauer, M. P., Krafft, T. D., Sesing-Wagenpfeil, A., & Zweig, K. (2023). *Quantitative study about the estimated impact of the AI Act.* arXiv. https://arxiv.org/abs/2304.06503

appliedAI. (2023, March). *AI Act: Risk classification of AI systems from a practical perspective. appliedAI Institute for Europe.* https://www.appliedai.de/en/insights/ai-act-risk-classification-of-ai-systems-from-a-practical-perspective/

OECD. (2024). *OECD artificial intelligence review of Germany.* OECD Publishing. https://doi.org/10.1787/609808d6-en

Hutchinson, P., Goll, F., & Mügge, C. (2024). *Generative AI in the European Startup Landscape 2024*. appliedAI Institute for Europe. https://www.appliedai-institute.de/en/hub/2024-generative-ai-study

Renda, A., Fanni, R., Laurer, M., Sipiczki, A., Yeung, T., Maridis, G., Fernandes, M., Endrodi, G., Milio, S., Devenyi, V., Georgiev, S., de Pierrefeu, G., & Arroyo, J. (2021). *Study to support an impact assessment of regulatory requirements for artificial intelligence in Europe: Final report.* European Commission, Directorate-General for Communications Networks, Content and Technology. https://doi.org/10.2759/523404

Fetic, L., Niemeyer, D., & Klein, T. (2025). *Wie können KI-Reallabore ihr Potenzial in der KI-Verordnung entfalten? Bedingungen für ein wirksames Instrument eines souveränen und lernfähigen KI-Ökosystems in Europa.* appliedAI Institute for Europe. https://www.appliedai-institute.de/media/downloads/Policy-Brief-AI-Regulatory-Sandbox.pdf?v=1764343275

Laurer, M., Renda, A., & Yeung, T. (2021, September 24). *Clarifying the costs for the EU's AI Act.* CEPS. https://www.ceps.eu/clarifying-the-costs-for-the-eus-ai-act/

Mueller, B. (2021, July). *How much will the Artificial Intelligence Act cost Europe?* Center for Data Innovation. https://datainnovation.org/2021/07/how-much-will-the-artificial-intelligence-act-cost-europe/

About the Authors

All authors contributed equally to this work.

Dr. Philip Hutchinson

Dr. Philip Hutchinson is a Senior AI Strategist at the non-profit appliedAI Institute for Europe. He also holds various expert positions, e.g. for the EuroHPC Joint Undertaking and the European Commission. Dr. Hutchinson previously worked as a management consultant at Ernst & Young, where he supported transformation projects in large companies. He holds a doctorate in innovation management from the University of Kiel with a research focus on the applicability of artificial intelligence in the innovation process.

Dr. Till Klein

Dr. Till Klein is the Head of AI Regulation at the non-profit appliedAI Institute for Europe gGmbH and the Project Manager of the Bavarian AI Act Accelerator. Till gained several years of industry experience in Regulatory Affairs for Medical Devices, as Lead Auditor for Quality Management Systems, and as Head of Quality Management in a Drone Company, which equips him with practical skills and perspectives that are instrumental for applying Trustworthy AI. He holds a Ph.D. in Business from the Centre for Transformative Innovation at the Swinburne University of Technology in Melbourne, Australia, and an undergraduate degree in Industrial Engineering.

Shahrukh Azhar Ahsan

Shahrukh Azhar Ahsan is an Applied AI Engineer experienced in deploying AI systems for private and government clients. He specializes in developing AI agents, RAG systems, and highly optimized inference pipelines for open-source LLMs and VLMs. Before focusing on AI, he worked with startups in the USA, UK and Canada developing DTC e-commerce platforms and speech-enabled web applications. He also co-founded an Ed-Tech AR startup for high school science laboratories. He holds an M.Sc. in Artificial Intelligence and Data Science from the Deggendorf Institute of Technology in Germany and a B.Sc. in Computer Science from the National University of Sciences & Technology (NUST) in Islamabad, Pakistan.

Akhil Deo

Akhil is a Senior AI Regulatory Affairs Expert at appliedAI Institute for Europe gGmbH, where he supports companies with enterprise AI governance and product policy. He previously worked as a communications specialist at the Future of Life Institute and as a research fellow at the Observer Research Foundation. Akhil holds a Master's degree in International Affairs from the Hertie School.

We as authors would like to express our sincere gratitude to Dr. Wolfgang Fischer and Alina Govoni for their valuable content guidance and insightful contributions to this publication. Their expertise and thoughtful input greatly enriched the quality and depth of our work.

Publisher
appliedAI Institute for Europe gGmbH
Freddie-Mercury-Straße 5
D-80797 München

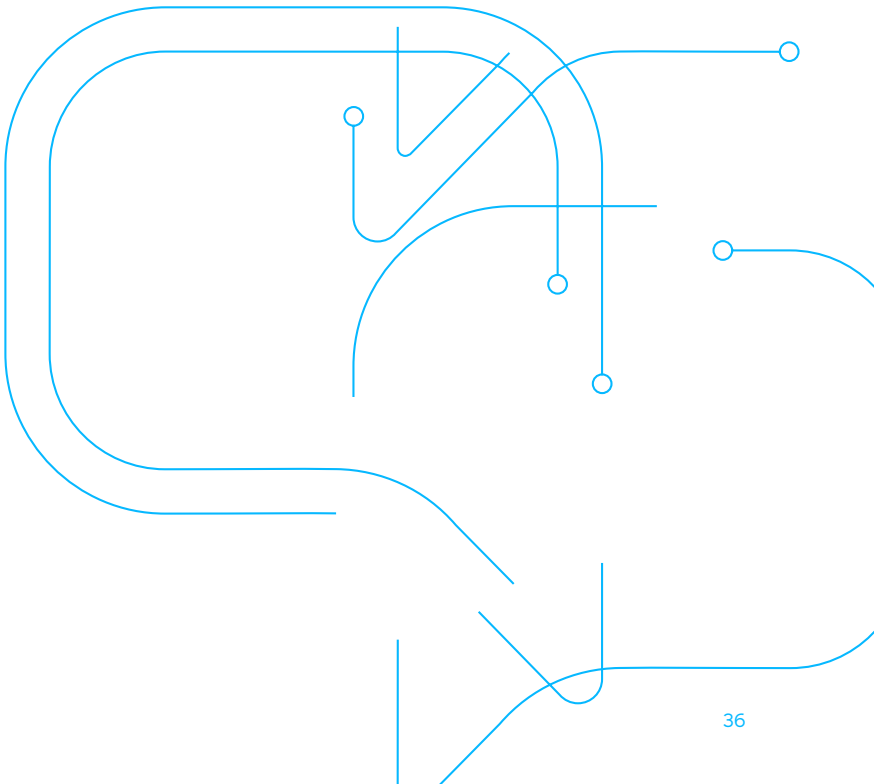

appliedAI Institute for Europe - About Us

The appliedAI Institute for Europe aims to strengthen the European AI ecosystem by engaging in research, developing knowledge around AI, providing trusted AI tools, and creating educational as well as interactive formats around high-quality AI content.

As a non-profit subsidiary of the appliedAI Initiative, the Institute was founded in Munich in 2022. The appliedAI Initiative itself is a joint venture of UnternehmerTUM and IPAI. The Institute is managed by Dr. Andreas Liebl.

The appliedAI Institute for Europe focuses on the people in Europe. It pursues the vision of shaping a common AI community and providing high-quality content in the age of AI for the entire society. By promoting trustworthy AI, the Institute accelerates the application of this technology and strengthens trust in AI solutions.

With a focus on research, knowledge development, research and the provision of trusted AI tools, the appliedAI Institute for Europe provides a valuable resource for companies, organizations, and individuals looking to expand their knowledge and skills in AI. Through educational and interaction formats, the Institute enables an intensive exchange of expertise and fosters collaboration between stakeholders from different fields.

The appliedAI Institute for Europe invites companies, organizations, startups, and AI enthusiasts to benefit from the Institute's diverse offerings and resources. The appliedAI Institute for Europe is supported by the KI-Stiftung Heilbronn gGmbH.

For more information, please visit www.appliedai-institute.de.

Gefördert durch
Bayerisches Staatsministerium
für Digitales

This study is part of the Bavarian AI Act Accelerator project.

The Bavarian AI Act Accelerator is a two-year project funded by the Bavarian State Ministry of Digital Affairs to support SMEs, start-ups, and the public sector in Bavaria in complying with the EU AI Act. Under the leadership of the appliedAI Institute for Europe and in collaboration with Ludwig Maximilian University, the Technical University of Munich, and the Technical University of Nuremberg, training, resources, and events are being offered. The project objectives include reducing compliance costs, shortening the time to compliance, and strengthening AI innovation. To achieve these objectives, the project is divided into five work packages: project management, research, education, tools and infrastructure, and community.

https://innovationsbeschleuniger.bayern

appliedAI institute for europe

Gefördert durch
Bayerisches Staatsministerium für Digitales

Bayerischer KI-Innovationsbeschleuniger

appliedAI Institute for Europe
Freddie-Mercury-Straße 5
80797 München
Germany
www.appliedai-institute.de