Crafting the AI Compliance Blueprint

Understanding and embedding AI regulations through a comprehensive AI Compliance approach

Fabian Meyer
Mauritz von Lenthe
Carola Bader
Muskaan Multani

July 2024
Blogpost

Key Facts

- **Rapid AI Advancement:** AI technologies are evolving at an unprecedented pace, with products like GPT-4, Gemini, and Claude demonstrating remarkable capabilities across various domains.
- **Complex Regulatory Landscape:** The global AI regulatory environment is fragmented, intricate and fluid, with the EU AI Act emerging as a significant framework that introduces a risk-based approach to AI governance; however, similar policy frameworks are also coming into force or in the process in various parts of the world like the US, Brazil, China and India.
- **Multifaceted Compliance Challenges:** Organizations face diverse AI compliance risks, including data privacy, governance accountability, and information security, each carrying potential legal and reputational implications.
- **Compliance Strategy:** To navigate this complex landscape, organizations must adopt a "compliance by design" approach, integrating regulatory considerations into their AI development lifecycle from the outset.
- **AI Compliance Blueprint:** A structured six-step methodology has been developed to guide organizations in embedding accountability, transparency, and responsible practices throughout their AI initiatives.
- **Cross-functional Governance:** Effective AI compliance necessitates the establishment of a robust governance framework, involving stakeholders bringing specialized expertise needed to address the complexity of modern AI systems, as well as from the entire AI value chain to ensure comprehensive oversight and risk management.
- **Success factors:** Success in AI compliance hinges on ongoing risk assessment, stakeholder engagement, robust data governance, and a commitment to continuous learning and adaptation to evolving regulatory standards.

Striking the right balance: with increasing regulatory pressure, organizations have to create a framework to ensure compliance in their AI initiatives

The remarkable progress in AI technologies like GPT-4, Gemini, and Claude showcasing impressive multimodal capabilities and human-level performance on several benchmarks is just one example of how artificial intelligence is transforming various industries and creating new opportunities. Gemini Ultra has become the first language model to achieve human-level performance on the Massive Multitask Language Understanding (MMLU) benchmark, with its performance improving by 15 percentage points over the past year. Additionally, GPT-4 has demonstrated a remarkable 0.96 mean win rate score on the comprehensive Holistic Evaluation of Language Models (HELM) benchmark, which includes MMLU among other evaluations.[1] However, AI as a broader field encompasses a wide range of technologies and use cases beyond just large language models (LLMs). While current AI systems still face challenges with factual accuracy, complex reasoning, and explainability, the rapid advancements that are being witnessing are indicative of the transformative potential AI holds across corporate organizations and multiple sectors.

Across industries, organizations are making substantial investments in AI, gradually moving beyond experimental stages to implement industry-ready solutions that deliver real business value. Driven by the promise for significant cost savings, operational advantages, and automation efficiencies the use cases observed vary from optimizing supply chains and enhancing customer experiences to streamlining financial processes and improving product development. This widespread adoption underscores AI's transition from a future prospect to a current, indispensable tool for competitive advantage in today's business landscape.

As AI becomes more pervasive and powerful, it also raises important regulatory and ethical questions that need to be addressed. A heated debate has emerged between those who prioritize innovation and those who advocate for stricter regulation to mitigate the potential risks and negative impacts of AI.

Proponents of innovation argue that excessive regulation could stifle the development and adoption of AI technologies, hindering their potential to drive economic growth, improve efficiency, and solve complex problems. They believe that industry self-regulation can effectively address AI-related concerns, arguing that overly prescriptive regulations may quickly become obsolete in this fast-evolving field.

Conversely, advocates for regulation stress the importance of establishing clear guidelines and safeguards to ensure AI systems are developed and deployed responsibly, ethically, and reliably. They emphasize potential risks such as privacy violations, discrimination, job displacement, and existential threats if AI is left unchecked. These supporters argue that a lack of oversight and accountability could lead to severe consequences for individuals, societies, and the environment.

This tension between innovation and regulation has led to a complex and dynamic regulatory landscape surrounding AI.

[1] https://aiindex.stanford.edu/wp-content/uploads/2024/04/HAI_2024_AI-Index-Report.pdf

In April 2021, the European Commission proposed the first ever legal framework on AI, which aims to ensure that AI is trustworthy, human-centric, and aligned with the EU's values and fundamental rights. The EU AI Act introduces a risk-based approach to regulation, with different levels of requirements and prohibitions depending on the potential impact of the AI system on human safety, dignity, and autonomy. The proposal also establishes a governance mechanism for the oversight and enforcement of the rules, as well as a coordinated plan for the development and uptake of AI in the EU. However, the EU is not the only actor in the global AI arena. Many other countries and regions have also developed or are developing their own policies, guidelines, standards, and strategies on AI, reflecting different perspectives and priorities.

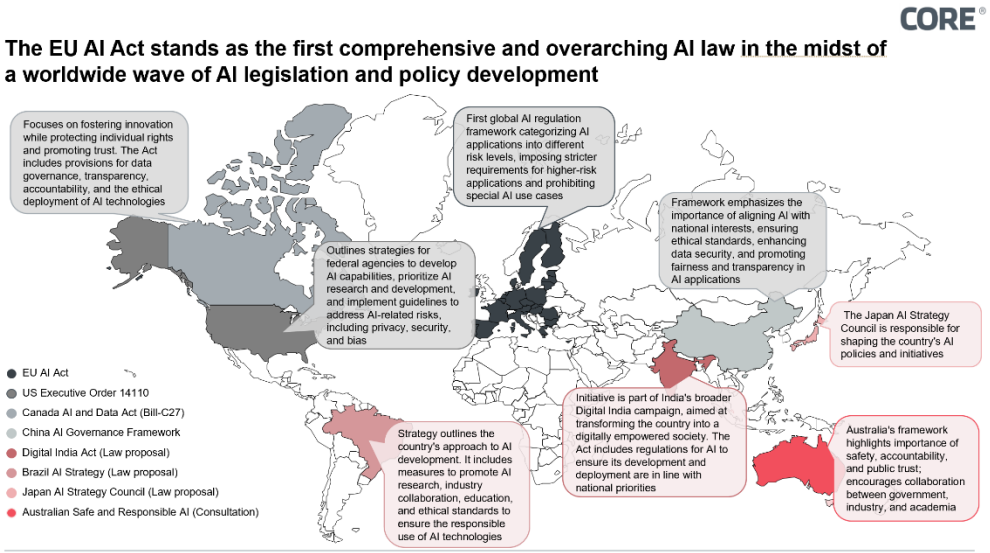

Figure 1: Global AI legislation and policies

For example, the US has issued several executive orders and national initiatives on AI, focusing on innovation, competitiveness, and national security. China has also adopted a comprehensive approach to AI, with ambitious goals to become a world leader in AI by 2030.[2] Singapore published their Model AI Governance Framework as a guidance to build a trusted ecosystem around AI.[3] And Japan has – in addition to its national AI strategy – published a set of further policies and guidelines such as the Social Principles of Human-Centric AI or the amendments to the copyright law to promote the use of data in machine learning by clarifying that downloading or processing data through the internet or other means to develop AI models is not an infringement of copyright.[4] Moreover, there are various multilateral and multi-stakeholder efforts to promote international cooperation and alignment on AI, such as the OECD Principles on AI, the G20 AI Principles, and the Global Partnership on AI.

This complex and dynamic regulatory environment poses significant challenges for organizations that want to adopt and deploy AI solutions. On the one hand, they need to keep up with the speed of technological change and leverage the potential benefits of AI for their business objectives. On the other hand, they need to comply with the relevant laws and regulations, as well as the ethical

[2] https://datagovhub.elliott.gwu.edu/china-ai-strategy/

[3] https://www.pdpc.gov.sg/help-and-resources/2020/01/model-ai-governance-framework

[4] https://www.csis.org/analysis/japans-approach-ai-regulation-and-its-impact-2023-g7-presidency

and societal expectations, that apply to their AI systems. Failing to do so can result in legal sanctions, reputational damage, and loss of trust among customers, partners, and regulators. For example, the EU AI Act stipulates that companies that violate the regulation can be fined up to 7% of their global annual turnover or €35 million per violation.

The rapid evolution of AI technology and the corresponding regulatory environment underscores a dynamic interdependency between technological advancements and legislative frameworks. Organizations are faced with increasing AI regulation and increased pressure to utilize AI at the same time. Therefore, it is essential for organizations to develop a robust AI compliance framework that can mitigate the risks and ensure the accountability, transparency, and fairness of their AI systems. Moreover, this framework should not be an afterthought or a reactive measure, but rather a proactive and strategic approach that is embedded in the design and development of the AI system from the outset – **compliance by design** is the keyword here. Crafting an effective AI compliance blueprint right from the beginning is a necessity, not a luxury, for any organization that wants to create a responsible and trustworthy AI future. To achieve this, a set of steps and best practices is suggested that can help organizations to implement a robust compliance management system along the entire AI value chain, involving all the relevant stakeholders and processes. While the blueprint is applicable to all sorts of organizations and regions, the focus of understanding the regulatory challenges and the corresponding solution will be on Europe in the following.

AI regulation in the EU: it’s not only about the EU AI Act

Navigating the Regulatory Challenges of AI Compliance

The regulatory challenges of AI compliance for organizations are multifaceted and complex, necessitating a comprehensive approach to managing the various risks associated with AI systems. These risks, which for example include bias, privacy concerns, accountability issues, and unintended consequences, can have significant legal and reputational implications for organizations. These key risks can be categorized into the following:

- Data Management and Privacy
- Governance and Accountability
- Information Security Risks

Data Management and Privacy includes the challenge of data quality and bias as well as complying with data protection regulations. AI systems can inadvertently perpetuate existing biases present in the data they are trained on, leading to unfair outcomes. For example, an AI system in hiring might favor certain demographics if training data isn't representative, leading to discriminatory practices and worsening social inequalities. Privacy is another major risk, as AI's use of personal data raises significant privacy issues. Organizations must comply with data protection laws like GDPR to avoid legal issues and maintain user trust. AI systems often require vast amounts of data to function effectively, and this data can include sensitive personal information. Establishing a strong Data Governance framework not only helps to ensure compliance with data protection regulations but also ensure data quality to mitigate the risk of bias by including for example statistical bias monitoring methodologies.

Governance and Accountability encompasses several challenges such as clear accountability, transparency, and ethical considerations as well as human oversight. Determining who is responsible for the actions and decisions made by AI systems can be complex, particularly when AI decisions lead to unintended consequences. Clear accountability frameworks are necessary to assign responsibility and ensure that there are mechanisms in place for redress. This involves defining the roles and responsibilities of AI developers, deployers, and users within an organization, ensuring that accountability is maintained throughout the AI lifecycle. Moreover, unintended consequences of AI systems can produce unexpected results with significant legal and reputational implications. For instance, a predictive policing AI might unfairly target specific communities, leading to public outcry and legal challenges. Beyond managing technical and legal risks, organizations must also address several ethical considerations. Transparency is crucial, as users and stakeholders need to understand how AI systems make decisions. Transparency helps build trust and ensures that AI systems are not perceived as "black boxes". Techniques such as explainable AI (XAI) aim to make AI decision-making processes more interpretable, providing explanations for their decisions and actions. Fairness in AI systems is another essential consideration. AI systems should be designed and trained to ensure they do not produce discriminatory outcomes. This involves careful consideration of the data used and the methods employed to mitigate bias. Fairness in AI includes ensuring equal treatment and opportunities across different demographic groups. Human oversight remains critical as it ensures that AI decisions align with ethical standards and societal values, allowing for intervention when necessary. This involves having mechanisms in place for humans to monitor, review, and override AI decisions, particularly in high-stakes applications such as healthcare and criminal justice. However, ensuring all those principles is faced with various challenges in practice, especially ensuring explainability and transparency. As seen with the growing investment in LLM development, the volume of data required to train billion parameter models has become so vast, that auditing that process has become intractable. The scale has reached a point where human attention is insufficient for conducting a singular audit, necessitating the reliance on tools and strategies to carry out these tasks. Implementing the requirements of legislation is one key step for this, especially the mandate for governance, thorough documentation and a robust data management process. In addition, joint efforts should be considered between companies. For example, in the US voluntary AI commitments are being fostered aimed at auditing and red-teaming each other, establishing information sharing etc.[5]

Finally, AI systems face considerable **Information Security Risks** that can impact their reliability and trustworthiness. A major challenge is their vulnerability to attacks of information exfiltration and the use of AI in augmenting believable phishing attacks at scale. For example, AI-driven chatbots can be manipulated to extract sensitive information from users by mimicking legitimate communication, making it difficult for individuals to distinguish between genuine and malicious interactions. This can lead to unauthorized access to confidential data and significant security breaches. Additionally, AI systems often handle large amounts of sensitive data, making them prime targets for data breaches. Protecting this data from unauthorized access and cyber-attacks

[5] https://www.whitehouse.gov/wp-content/uploads/2023/09/Voluntary-AI-Commitments-September-2023.pdf

is crucial but difficult. Furthermore, the complexity of AI algorithms can make it challenging to detect and fix security vulnerabilities quickly.

While organizations should address all these potential pitfalls proactively, they are more and more also being forced to implement such mitigation measures as regulators are becoming active.

Regulatory Frameworks Addressing AI Risks

To navigate these challenges, regulators have stepped in to provide frameworks that guide the compliant use of AI across its lifecycle. In the European Union, the most comprehensive of these frameworks is the EU AI Act, which provides a risk-based approach to the production, use, and distribution of AI with a strong emphasis on ethics, information security, and data protection. However, the regulatory landscape is not limited to the EU AI Act. Several additional frameworks contribute to the overall AI regulation, addressing different stages of the AI process – input, generating, and output – and distinguishing between EU-level and national-level regulations:

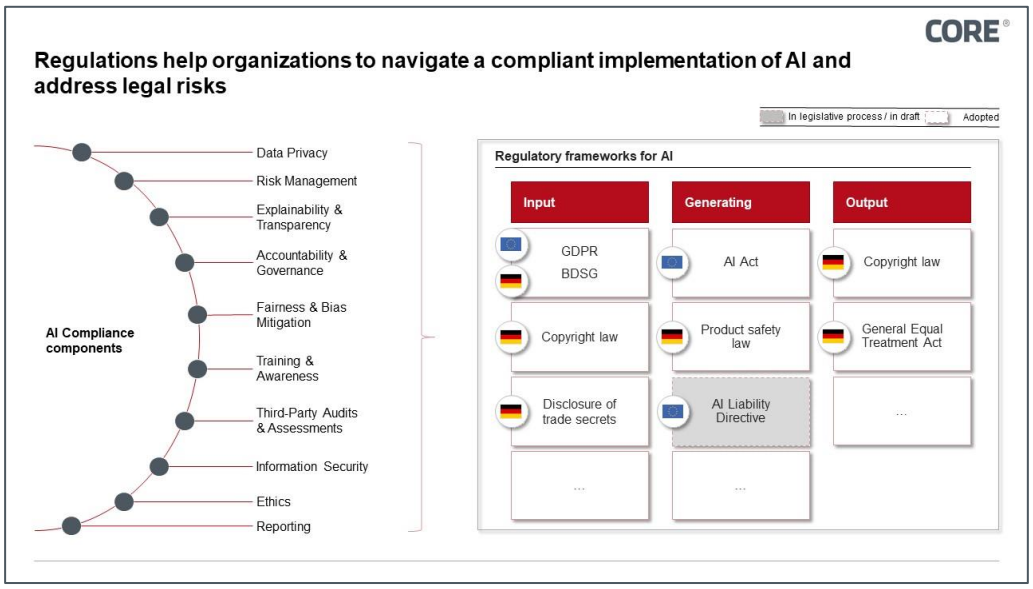

Figure 2: AI regulatory frameworks in Germany

Input: Legal risks at this stage primarily involve data protection and privacy concerns. The GDPR (General Data Protection Regulation) and – for example in Germany – the BDSG (Bundesdaten-schutzgesetz) are pivotal in regulating the collection and use of personal data. The GDPR ensures that personal data is handled with care, requiring informed consent, legal permissions for processing, and regular data protection impact assessments. The BDSG complements these regulations in Germany, adding specific provisions for sensitive data. These laws mandate that AI systems manage personal data transparently and securely, emphasizing the protection of individual rights. Compliance with these regulations is essential for developing trustworthy and legally compliant AI technologies.

Additionally, the use of third-party intellectual property rights, especially copyright laws, is critical to ensure that AI systems do not infringe on existing intellectual property. The German Copyright Act (Urheberrechtsgesetz) is particularly relevant here.

Moreover, the disclosure of trade secrets is governed by the German Trade Secrets Act (Gesetz zum Schutz von Geschäftsgeheimnissen), which ensures that proprietary information is adequately protected.

- **Generating:** The generation phase of AI involves several regulatory frameworks to ensure the legality and safety of AI systems. The EU AI Act is central to regulating the use and development of AI technologies, emphasizing compliance with ethical standards and mitigating potential risks.

Figure 3: Overview on the risk-based approach and roles in the EU AI Act

It introduces a risk-based approach to categorizing AI systems, focusing on the potential impact on individuals and society. AI systems are classified into three categories – unacceptable risk, high risk, and systemic risk:

- **Unacceptable Risk:** AI systems that pose a clear threat to safety, livelihoods, and rights are prohibited. This includes AI applications like social scoring by governments and real-time biometric identification in public spaces without judicial oversight.
- **High Risk:** These AI systems are subject to stringent requirements due to their significant impact on individuals' rights and safety. High-risk AI includes technologies used in critical infrastructure, education, employment, essential public services, law enforcement, and migration management. The requirements for high-risk systems are a risk management system, data and data governance, technical documentation, record keeping, transparency and information to users, human oversight, accuracy, robustness, and cybersecurity.
- **General-Purpose AI Systems (GPAIs) with Systemic Risks**: These systems, which can broadly impact public safety, democracy, and the rule of law, require comprehensive oversight. They include AI technologies like generative models and predictive policing systems. GPAIs must adhere to strict governance frameworks and regular evaluations to mitigate potential widespread negative impacts.

The EU AI Act outlines specific compliance obligations that vary according to the role of the entity involved. There are four key roles: providers, importers, distributors, and deployers, each with distinct responsibilities in the AI ecosystem

- Providers are entities that develop and place AI systems on the market, ensuring these systems meet all regulatory standards, including risk management, data governance, and transparency
- Importers are responsible for verifying that the AI systems they bring into the market from outside the EU comply with the regulations
- Distributors supply AI products within the market and must guarantee these products adhere to regulations and maintain detailed records
- Deployers, or end-users, implement AI systems in real-world applications and need to ensure the ethical and legal use of these systems, including establishing oversight mechanisms.

These obligations differ between high-risk AI systems and GPAIs with systemic risks, ensuring tailored regulatory oversight for each AI application category. Additionally, there are overarching duties such as adhering to ethical principles in the use of AI and ensuring that staff and users have sufficient AI literacy. In this context, providers and deployers must ensure their staff and users of AI systems have adequate technical knowledge, experience, education, and training, considering the context and the people affected by the AI systems.

Product safety laws, such as the German Product Safety Act (Produktsicherheitsgesetz), ensure that AI products meet the required safety standards to prevent harm to users. Therefore, these laws mandate rigorous testing and compliance checks before AI systems can be deployed in the market.

Additionally, the EU AI Liability Directive, which is still in the draft stage, aims to streamline liability assignment, establish a rebuttable presumption of causality, and ease the proof burden for affected parties. This directive is designed to complement the EU AI Act by addressing legal complexities and ensuring that organizations can effectively manage risks associated with AI. While it is not yet finalized, the directive underscores the importance of having a robust risk management framework to handle liability issues related to AI systems.

- **Output**: The final stage of the AI process focuses on the outcomes and their impact. Legal risks here include issues of bias and misinformation, which can significantly affect public trust and fairness. The German General Equal Treatment Act (Allgemeines Gleichbehandlungsgesetz) is relevant in addressing potential discriminatory outcomes of AI systems. Additionally, infringement of third-party intellectual property rights continues to be a concern, necessitating compliance with copyright protection laws. Moreover, the use of AI-generated content is governed by licensing restrictions, ensuring that the output adheres to legal standards and respects intellectual property rights.

In summary, these frameworks ensure that AI technologies are developed and used responsibly, promoting transparency, fairness, and human oversight. While the EU AI Act and the EU AI Liability Directive are explicitly designed for AI, other regulations such as the GDPR apply in specific AI contexts, ensuring comprehensive oversight. As regulations continue to evolve, organizations must stay informed and compliant, leveraging these frameworks to build trustworthy

and ethical AI systems. By understanding and adhering to these regulations, organizations can effectively manage the risks associated with AI and foster an environment of trust and accountability.

From Challenge to Compliance: the AI Compliance Blueprint

The regulatory requirements for AI systems translate into various components that organizations must address to ensure compliance. To identify these components, a six-step approach is suggested, **the AI Compliance Blueprint**. Implementing a robust AI Compliance Blueprint is crucial for organizations aiming to develop and deploy AI systems responsibly and ethically. 1.

Figure 4: AI Compliance Blueprint

1. Establishing clear guiding principles

Core values and principles need to be defined that will inform and shape the AI compliance strategy and culture. These principles should reflect the ethical and societal expectations and standards that apply to the organization and its AI systems, as well as the organizational goals and vision – thus overall alignment with the AI strategy is needed. This also includes decision-making and alignment on where and how AI will be used and communicated with stakeholders. In general, a mix of so called hard and soft guiding principles is suggested for a comprehensive set of guiding principles suited to guide everyone in the organization.

Hard guiding principles emerge from regulatory requirements such as transparency, risk management and monitoring requirements. They should be aligned with the existing or emerging regulatory frameworks and industry standards e.g. the EU AI Act or the OECD Principles on AI.

Soft guiding principles on the other side are derived from organizational values, ethical considerations, and societal expectations and similar not directly coded in legislation. These principles could include for example sustainability, ethical conduct and user-centricity.

2. Establishing an AI governance framework

The second step is to set up the organizational structures and roles that will oversee and manage the AI compliance processes and activities. Effective governance is the backbone of a robust AI compliance program. It involves establishing an AI Governance Body with clearly defined policies, procedures, roles, and responsibilities. This body plays a crucial role in operationalizing legal requirements and ensuring that the necessary governance structure is in place to navigate the complexities of AI regulation and compliance. The AI Governance Body should be composed of cross-functional representatives from various departments and disciplines within the organization. This diversity of perspectives ensures that all aspects of AI development, deployment, and compliance are considered, fostering a comprehensive and well-rounded approach. One of the primary responsibilities of the AI Governance Body is to develop and implement policies and procedures that align with the organization's responsible AI principles and compliance objectives. Defining roles and responsibilities is another critical task for the AI Governance Body. This includes identifying key stakeholders, subject matter experts, and decision-makers, and outlining their respective duties and accountabilities within the AI compliance framework. In addition, required reporting structures need to be put in place. The reporting requirements differ based on the role of an organization but alignment with the national supervisory authority should be ensured and maintained at all times.

Senior management support is vital for the success of the AI Governance Body and the overall compliance program. Executive leadership must champion the importance of responsible AI practices and compliance, allocating the necessary resources and fostering a culture of accountability and continuous improvement. By establishing a robust governance structure, organizations can ensure that their AI initiatives are developed and deployed in a responsible, ethical, and compliant manner, mitigating risks and fostering trust among stakeholders and the broader public.

3. AI Maturity Assessment

The third step is to assess the current state and readiness of the organization for AI compliance, as well as the maturity and quality of the AI systems and practices. This involves setting up an AI inventory and then conducting a comprehensive and systematic analysis of the strengths, weaknesses, opportunities, and threats (SWOT) related to AI compliance, as well as identifying the gaps and areas for improvement. In detail, organizations should conduct:

- Assessments of documentation and documents on infrastructure, platforms and applications
- Interviews with relevant key stakeholders on current state of AI use cases, *a) pursued* and *b) achieved* objectives and estimated potential
- Evaluations of organisational (AI) capabilities, along existing roles and resources

This involves evaluating the maturity and quality of the AI systems and practices, using relevant criteria and indicators, such as data quality, model accuracy, robustness, explainability, scalability, and usability. The results of the assessment will provide a baseline and benchmark for

designing and implementing the AI compliance processes and activities, as well as measuring and improving their effectiveness and efficiency.

4. Mapping exercise

The step is to ensure that the AI systems and practices comply with the applicable laws and regulations, as well as the relevant industry standards and best practices. This includes identifying and mapping the legal and regulatory requirements and obligations that apply to the AI systems and initiative. Additionally, if an organization adheres to voluntary standards such as the ISO/IEC 27001 for information security management or the ISO/IEC 38500 for IT governance, these can be in scope of this mapping exercise as well to ensure alignment. Furthermore, voluntary standards also can offer guidance for the mapping e.g. the NIST AI 60-1 for GenAI systems categorization.[6]

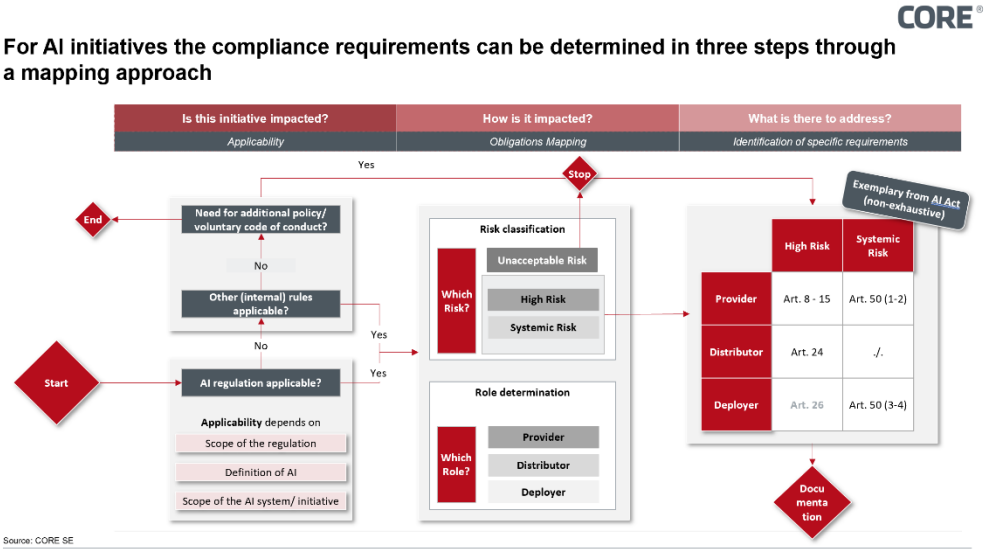

Figure 5: Process view on mapping between AI regulation and AI initiatives

The mapping approach starts by assessing whether the AI regulation is applicable, considering factors such as the scope of the regulation, the definition of AI, and the scope of the AI system or initiative. If the regulation is applicable, the process continues by evaluating the need for additional policies or voluntary codes of conduct, as well as considering other internal rules that may be relevant. The flowchart then guides the user through a risk classification step, where the risk is categorized as unacceptable, high, or systemic.

Based on the identified risk level, the next step is to determine the specific role (provider, distributor, or deployer) and map the corresponding obligations and requirements outlined in the regulation. For example, for high-risk systems a list of requirements is identified depending on the role of the organization reaching from organizational requirements such as a governance structure and informing employees to technical requirements like quality management systems and documentation keeping to reporting and cooperation requirements with the respective authorities. Furthermore, this includes implementing and maintaining the necessary

[6] https://airc.nist.gov/docs/NIST.AI.600-1.GenAI-Profile.ipd.pdf

documentation, certification, and auditing mechanisms to demonstrate and verify compliance with the legal and regulatory frameworks and industry standards.

Legal obligations for high-risk AI systems under the EU AI Act differ depending on the role

	Legal obligations (incl. main referal articles)	Providers (Art. 16)	Importers (Art. 23)	Distributors (Art. 24)	Deployers (Art. 26)
1	Establishment of a **risk management system** (Art. 9)	✓			
2	Requirements regarding **training, validation and testing data** (Art. 10)	✓			
3	Technical **documentation** (Art. 11)	✓	✓		
4	**Record keeping** (Art. 12)	✓			
5	**Transparency** and provision of information to deployers (Art. 13)	✓			
6	Human **oversight** (Art. 14, 26(2))	✓			✓
7	**Accuracy, robustness** and **cybersecurity** (Art. 15)	✓			
8	**Quality** management system (Art. 17)	✓			
9	**Documentation** keeping (Art. 18, 23(5))	✓			✓
10	**Corrective actions** and provision of **information** (Art. 20, 23(2), 24(2)(4))	✓	✓		
11	**Conformity** assessment (Art. 43)	✓	✓		
12	**Registration** obligation (Art. 49)	✓			✓
13	**Cooperation** with **competent authorities** (Art. 21, 23(6), 24(5)(6), 26(12))	✓	✓	✓	✓
14	**EU declaration** of conformity (Art. 47)	✓	✓	✓	
15	Affix **CE marking** (Art. 48)	✓	✓	✓	
16	Indication of **name/address** (Art. 16(b), 23(3))	✓	✓		✓
17	**Comply with instructions** for use (Art. 26(1))				✓
18	Consider **relevance of input data** (Art. 26(4))				✓
19	**Monitor operation** of the system (Art. 26(5))				✓
20	**Record keeping** for **automatically** generated **logs** (Art. 19, 26(6))	✓			✓
21	Execution of **data protection impact assessment** (Art. 26(9))				✓
22	Information of **employers/natural persons** (Art. 26(7)(11))				✓

Figure 6: Legal obligations for high-risk AI systems depending on the role

5. Gap Analysis

Once the applicable requirements have been identified, conducting a thorough gap analysis is the next critical step. This process involves a comprehensive assessment of the organization's current state regarding AI rules, including the existing compliance setup, rules, processes, and practices related to AI development and deployment. The goal is to identify all gaps, no matter how small or seemingly insignificant, as even minor non-compliance issues can have significant consequences.

To begin, a cross-functional team that includes representatives from various departments such as legal, IT, data science, HR, and operations should be engaged. This is normally implemented through having a core project team but ensuring the support of SMEs from different departments if required. This ensures a holistic view of the organization's current practices and helps in identifying gaps from different perspectives.

Next, all relevant documentation, including current policies, procedures, workflows, training materials, and compliance records should be gathered as a preparation measure. This provides a baseline for comparison against regulatory requirements.

The gap analysis consists of two essential parts:

6. Regulatory requirements vs. Policies and Governance

This part of the gap analysis involves comparing the specific regulatory requirements with the organization's existing policies and governance structures. The goal is to assess whether the policies, procedures, and governance frameworks in place adequately address the compliance obligations set forth by relevant AI regulations. As a result, an organization is able to identify

whether all required policies are in place (i) and if those policies have the minimum required content (ii).

7. Regulatory requirements vs. operational implementation/ actual state

The second part of the gap analysis focuses on the operational implementation of policies and the actual state of AI practices within the organization. This involves examining whether the day-to-day operations and the actual use of AI systems align with both regulatory requirements and the organization's own policies. Engaging with key stakeholders through interviews and surveys to gain insights into actual practices and procedures, ensuring that documented processes are being followed in practice and uncovering any undocumented practices is key here. This analysis ensures that the actual implementation of AI systems and practices within the organization adheres to compliance standards, addressing any discrepancies between written policies and actual operations. This helps in maintaining consistent and effective compliance across all levels of the organization.

Once the gaps have been identified, the next step is to define clear and actionable measures to close them. This may involve revising existing policies and procedures, implementing new processes or controls, providing additional training to personnel, or allocating resources to address specific compliance requirements.

8. Implementation

After identifying the gaps and defining the action items, the last phase of the AI Compliance Blueprint is execution – making the necessary changes to achieve the desired target state and setting up the needed structures for continuous compliance. This requires sufficient implementation time as – depending on the results of the gap analysis and the defined measures – several efforts within the organization can be required.

Proving compliance is a vital aspect of this phase. It requires to set up mechanisms for documenting and showing the organization's compliance with the applicable AI regulations and frameworks. This may involve keeping detailed records, performing internal audits, or hiring external auditors to verify compliance efforts. As organizations shift towards AI-centric operating models, the generation of compliance-associated artifacts should become and integrated part of product development. Just as performance testing is part of software development, bias testing and other Responsible AI practices, can be built into the software development lifecycle and active monitoring workflows by design. This also includes making sure that the processes are in place to provide the requested information by the national supervisory authority or EU AI Office as well as have processes in place for reporting of incidents and breaches.

Tracking progress is also important. As regulations and best practices change, your AI compliance program must change accordingly. Awareness regarding changes in the regulatory environment and industry standards has to be maintained as well as an approach to update processes and controls to keep compliant.

Final remarks

In conclusion, the rapidly evolving landscape of AI regulation, particularly in the European Union, presents both challenges and opportunities for organizations. To navigate this complex environment and leverage AI's potential responsibly, organizations should focus on three key areas:

1) **Implementation of the AI Compliance Framework**
    - Establish guiding principles and an adoption framework
    - Implement continuous risk assessment and mitigation processes adopted throughout the AI development lifecycle
    - Ensure transparency and explainability in AI decision-making

2) **Prepare the technical foundation**
    - Implement robust data management and product-centric AI approach
    - Build fit-to-purpose organisational and technological infrastructure
    - Develop AI Governance System, leveraging learnings from Data Governance system where possible

3) **Create organisational and cultural readiness**
    - Provide comprehensive training across all organizational levels
    - Encourage collaboration and knowledge sharing
    - Provide secure environment to foster experimentation and innovation

Figures

Fabian Meyer is Managing Partner of COREconsulting with a focus on the internationalisation of CORE's services. In the client dimension, he is responsible for the implementation of complex IT projects with a focus on mergers & acquisitions, payments and transaction banking. He already gained experience as a management consultant and founder during his business studies, which he completed with a Master's degree in Mannheim. He has several years of consulting experience in the technology sector.

Mail: fabian_meyer@epam.com

Mauritz von Lenthe is a Senior Transformation Manager at CORE with a focus on Information Security & Compliance, especially in Data Management & AI. His versatile education from business administration to industrial engineering to data scientist gives him a unique perspective on technology-driven processes. As an experienced project manager and AI specialist, he develops innovative solutions that combine security, compliance and advanced data utilisation.

Mail: mauritz_lenthe@epam.com

Carola Bader is an Expert Associate at CORE and has extensive experience in the areas of compliance and risk management in the financial industry. She specialises in advising financial services providers on implementing effective compliance strategies, meeting regulatory requirements and combating financial crime.

Mail: carola_bader@epam.com

Muskaan Multani is a Transformation Associate at CORE. With her expertise and experience in information security & compliance, she focuses on supporting compliance projects, primarily in the financial industry, from strategic conception to implementation. With her educational background in business law and strategy & consulting, she brings an unique approach regarding regulatory compliance to organizational challenges. She is responsible for supporting project teams and clients in business-critical technology transformations.

Mail: Muskaan_multani@epam.com

COREtransform GmbH
Kurfürstendamm 194
10707 Berlin | Germany
https://core.se/
Phone: +49 30 263 440 20
sharedcoreoffice@epam.com

COREtransform GmbH
Limmatquai 1
8001 Zürich | Helvetia
https://core.se/
Phone: +41 44 261 0143
sharedcoreoffice@epam.com

COREtransform Ltd.
9 Devonshire Square, 5th Floor
London EC2 4YF
Great Britain
https://core.se/
Phone: +44 20 328 563 61
sharedcoreoffice@epam.com

