

## AI Governance for Third-Party Models: The Compliance Blind Spot in Vendor AI

Amit Awasthi

Dr. Bhimrao Ambedkar University, India

---

### ARTICLE INFO

Received: 22 Dec 2025

Revised: 08 Feb 2026

Accepted: 16 Feb 2026

### ABSTRACT

As organizations rapidly deploy AI to govern many core business processes, third-party and vendor AI models are increasingly used in healthcare, financial services, manufacturing, and public sectors. While organizations invest large budgets in governing their internally developed AI, vendor AI often lacks any governance, creating a compliance blind spot which can result in many risks and violations. This governance gap is particularly critical from a cybersecurity perspective, as vendor AI systems introduce unique attack surfaces including adversarial manipulation vulnerabilities, data exfiltration risks through model APIs, and supply chain security weaknesses that traditional security controls cannot adequately address. Regulatory regimes fully assign liability to organizations deploying AI, regardless of whether the organization developed the AI in-house, or purchased it. Because organizations lack technical access to the underlying components and external behavior of vendors' AI, they cannot scrutinize vendor AI for security vulnerabilities, model poisoning attempts, or backdoor exploits. This creates an intrinsic paradox: Organizations fully own responsibility for systems that they cannot fully control, audit for security flaws, nor understand. We propose a Third-Party AI Governance Framework based on systematic vendor AI classification based on decision criticality and regulatory ramifications, shared accountability architectures based on vendor-client responsibility, audit-ready transparency standards that require full documentation of all model attributes including security testing results, and continual oversight mechanisms beyond procurement-time evaluation. Organizations that cannot establish sufficient control over vendor AI security posture are likely to experience regulatory enforcement actions, security breaches, lawsuits, and loss of stakeholder trust. By contrast, showing and proving control of vendor AI through third-party AI governance may allow an organization to survive and even give it a competitive edge.

**Keywords:** Vendor AI Governance, Third-Party Model Accountability, Algorithmic Transparency, AI Compliance Frameworks, Enterprise AI Risk Management

---

### 1. Introduction

AI is permeating all aspects of organizations' daily operations and is transforming decision processes in all major business areas. AI technologies and applications are being used in core systems for healthcare, financial services, manufacturing, transportation and public service provision. This common AI consumption, and the accompanying step change in capabilities, is a major shift as AI transitions from an optional enhancement to a core business capability that determines competitiveness and efficiency.

The AI ecosystem has been shifting from organizations building and deploying enterprise AI models in house to AI-enabled consumption-oriented architectures, as evidenced by the rise in AI services through third-party software platforms, cloud machine-learning services and vendor-controlled algorithmic systems being embedded in applications [2]. With clinical decision support, machine learning principles are used to inform health care professionals about diagnostic and treatment recommendations. Third-party fraud detection systems are used by financial services companies to monitor and identify fraudulent transactions in real time, while third-party AI-improved customer analytics systems provide insights into

customer behavior and future preferences. In manufacturing environments, vendor-supplied predictive maintenance systems analyze equipment sensor data to forecast failures, AI-powered quality control systems detect production defects in real-time, and autonomous supply chain optimization platforms coordinate inventory, logistics, and production scheduling across global operations. Vendor-supplied autonomous optimization systems continuously optimize supply chains, energy grids, and factory conditions across plants without requiring human operators.

From a cybersecurity perspective, this proliferation of vendor AI introduces critical security challenges that extend beyond traditional software vulnerabilities. Vendor AI systems present unique attack vectors including adversarial examples that manipulate model predictions, model inversion attacks that extract sensitive training data, and supply chain compromises where malicious actors inject backdoors during model development or distribution. Organizations using vendor AI typically have low visibility into the vendor's security testing program, vulnerability disclosure program, or incident response plan, creating potential blind spots in the threat landscape.

However, reliance on externally provided AI functions gives rise to a governing asymmetry within the institutional AI portfolio. Internal AI development projects are usually subject to more strict governance frameworks including model testing, performance evaluation, security assessments and compliance verification protocols prior to production deployment [1]. For effective AI risk management, in-house created and managed algorithms include systems for thorough documentation, testing, security validation, and accountability that mirror enterprise risk management processes. In contrast, externally developed AI systems are often acquired through less technical organizational procurement processes that lack the same levels of oversight, security scrutiny, and rigorous control.

The governance gap is the disparity between regulatory accountability and operational control. These frameworks place accountability for the results of the use of AI systems in the hands of the organizations that implemented the systems, irrespective of where the models are sourced from [2]. Where AI systems produce biased outputs, harmful decisions, security vulnerabilities, or non-compliance with data protection and privacy legislation, the implementing organization will be held responsible even when organizations are not developing an underlying AI model but are buying and integrating an AI system as part of a larger software product. Either way, organizations typically lack the technical means to understand, document, audit, assess security posture, and provide oversight of the behavior of AI systems provided by vendors in order to meet compliance. This creates a compliance dilemma where organizations are responsible for the outcomes and security of AI systems they are not in control of and for which they cannot provide meaningful oversight.

## 2. The Transformation of Enterprise AI Adoption

The enterprise AI ecosystem moved from ownership to consumption. Within a year of generative AI's mass availability, 33% of respondents were using generative AI at least weekly within one business function [3]. The rapid adoptions have not merely been experimentation: among adopters of AI, 60% report using generative AI, signaling lasting organizational change in how capabilities are procured and used [3]. NIST characterizes those systems as billion-scale or trillion-scale and having billions or trillions of decision points; thus, their failure modes and emergent properties are difficult to anticipate in an organizational context [4].

The push for AI standards is part of a global trend towards establishing ethical guardrails on AI. To date, 47 countries have committed to the OECD AI Principles, the first intergovernmental AI standard [4]. In total, these principles urge AI actors to respect human rights, democracy, and rule of law throughout the different stages of the life cycle of AI systems defined by NIST: Plan and Design; Collect and Process Data; Operate and Monitor [4]. The transformation will focus on multiple high-value domains.

Marketing and Sales: The most widely used application, organizations are engaging in content creation and personalized consumer interactions through vendor models [3].

- Product and Service Development: External models are used to accelerate the research, development and design of a company [3].
- Service Operations: Vendor-managed tools offer customer support and automated troubleshooting [3].
- Manufacturing Operations: Third-party AI systems enable predictive maintenance to minimize equipment downtime, computer vision models for automated quality inspection and defect detection, demand forecasting algorithms for production planning, and robotics control systems for automated assembly lines. Manufacturing organizations increasingly rely on vendor AI for process optimization, energy consumption reduction, and real-time production scheduling across distributed facilities.

Though adoption of generative AI has occurred at a very rapid rate, some governance levers have yet to be fully engaged. For example, while Robustness, Security, and Safety emerged as trustworthy AI values in the OECD principles [4], only 21% of AI-using organizations report developing policies for employee use of external tools [3]. And only a third (32%) of organizations surveyed report that they reduce inaccuracy, the most commonly reported risk of third-party generative models [3], despite the emphasis placed by NIST on validity and reliability as core attributes of reliable systems [4].

The economic impact of generative AI alone is estimated to be between 2.6 trillion and 4.4 trillion dollars per year. In knowledge-intensive sectors like banking, pharmaceuticals, and advanced manufacturing, vendors using this technology are likely to derive 5% of the total industry revenue in value. In contrast, such gains in technology sectors could potentially reach 9% [3]. OECD criteria for enabling ecosystems with Transparency and Explainability information can provide an essential basis for sustaining growth enabled by persistent public trust [4]. NIST makes GOVERN an overarching requirement to inform all other parts of AI risk management [4].

### 3. Defining the Compliance Blind Spot

#### 3.1 The Accountability-Control Paradox

Regulators have also stressed the deployer responsibility principle. Organizations deploying AI systems are responsible for use of AI systems, whether developed in-house or procured from external sources. Deployers must ensure that systems deliver benefits safely, fairly, securely, and in compliance with the law [5].

A vendor's AI presents fundamental challenges for organizations' accountability, as its models are less auditable than internal models when the latter's training data, architectures, security testing results, and validation results are available. Vendor AI is a black box with little operational transparency and even less security transparency [6]. Organizations cannot audit algorithms, examine training data, evaluate adversarial robustness, or view security assessments and performance metrics stratified by demographics and deployment contexts.

From a security standpoint, this impracticability creates a challenge as organizations are unable to verify if the AI models provided by the vendor were tested against adversarial attacks, whether the training data was sanitized to prevent data poisoning, or if the model serving infrastructure is secured. As penetration testing, security audits and vulnerability assessments are rarely possible with vendor AI systems, organizations are exposed to risks that they cannot detect, measure or reduce. But there is a paradox: legal regimes place the burden of full liability on organizations that cannot reach this level of control. Legal obligations do not map easily onto technical controls. So organizations are tasked with meeting legal obligations in the absence of knowledge about the systems and how they function [5].

**3.2 Mechanisms of Inherited Risk**

For vendor AIs, the latent risk mainly comes from generally poorly understood factors related to different aspects of inherited risk from a security perspective:

Training datasets can be biased, e.g., when they are derived from a training set that does not include underrepresented or historically disadvantaged groups and/or biased by labeling assumptions made when the dataset is created [6]. Further, the dataset can also be biased if there is a backdoor, meaning they have been deliberately poisoned. Clients will take on the risk of the vendor's security posture on data that was used to train the LLM (while having no visibility to provenance).

Model security vulnerabilities: Security vulnerabilities in model serving infrastructure can be used by attacking models. These vulnerabilities can include unsecured vendor-serving of models, lack of defenses against adversarial examples that induce a model to output a different prediction than the expected one, lack of input validation during model serving that can be exploited by an injection attack, and data exfiltration via model inversion attacks or via membership inference attacks [5]. Vendor-serving systems are common, and the same attack can be used against multiple organizations.

Supply Chain Risk: The AI model supply chain is similar to the software supply chain but with additional complexity. Compromised development environments, malicious code injection during model training, tampering with model weights during distribution, and unauthorized modifications to deployed models all represent supply chain attack vectors that organizations cannot detect without deep technical access to vendor systems.

Systemic risk amplification occurs when risk related to individual vendors spreads among sectors. This risk is especially high in single vendor models due to thousands of organizations relying on a single vendor whose foundational vulnerability affects all [6]. A security flaw in a widely-deployed vendor AI model could create cascading failures across industries, from financial fraud detection systems to manufacturing quality control to healthcare diagnostics.

**3.3 Compliance Assumption Failures**

Organizations often operate under the assumption that having a procurement department that treats certifications (e.g., SOC 2 reports, ISO certifications, etc.) as proxies for compliance is sufficient. However, these generic certifications address how an organization implements systems and processes, not how an AI model behaves or how secure it is against AI-specific threats. Traditional security certifications do not cover adversarial robustness testing, model poisoning defenses, or AI-specific vulnerability assessments.

As a result of increased regulatory scrutiny on controls demonstrated rather than assumed, such frameworks require an organization to have commitment to controls evidenced through validation of controls through testing, operational metrics or activities. Assertions or contractual representations are insufficient substitutes [5]. Security attestations without technical validation evidence leave organizations exposed to both compliance violations and actual security incidents.

<b>Risk Category</b>	<b>Source</b>	<b>Impact Scope</b>
Training data bias	Upstream datasets	Discriminatory outcomes across demographics
Training data poisoning	Malicious dataset injection	Backdoor triggers and model manipulation
Security vulnerabilities	Vendor infrastructure	Data exfiltration and adversarial attacks
Adversarial robustness gaps	Insufficient security testing	Prediction manipulation and system compromise
Supply chain compromise	Development/distribution tampering	Backdoored models and unauthorized behavior

Unannounced model updates	Vendor changes	Behavioral modifications without validation
Jurisdictional misalignment	Regulatory differences	Compliance gaps across markets
Systemic risk amplification	Single vendor deployment	Sector-wide failure propagation

Table 1: Inherited Risk Categories in Vendor AI Deployment [5, 6]

#### 4. Gaps in Current AI Governance Paradigms

Governing AI today means supervising internally developed AI systems. The AI governance models adopted by modern organizations are development-focused, helping govern their internal AI model lifecycle processes, i.e., the organization sets policies for AI development from conception and initial development to production deployment and deprecation [7]. Organizations spend considerable effort documenting the entire internal model development process for how the features were engineered, selected, tuned, and validated.

Another important area of focus for contemporary AI governance frameworks is internal data governance frameworks developed by organizations to guide the organization's data assets feeding into AI systems. These include data quality assurance tasks such as validation rules and data cleansing, data lineage for tracking the source and transformation of datasets, access controls limiting exposure to sensitive data used in training, and the protection of personally identifiable information (PII) during the AI model development life cycle [8]. Internal data governance is vital since the model quality depends heavily on the quality of training data. All models trained on biased, incomplete or low-quality data will result in bad models.

Governance frameworks are thorough but they also suffer from critical shortcomings across the board of commercial AI systems. No governance framework has provisions for communicating with third-party model transparency mechanisms, and organizations have no formalized processes or tools to request, analyze, and review the AI systems used by vendors to perform services [8]. However, internal models' full visibility into their own training data, architectures, and performance characteristics offers little guidance on what organizations should expect from vendors, how to evaluate vendors' level of transparency, or even how organizations should respond in cases in which vendors refuse to share information regarding how to govern their models.

A lack of shared accountability frameworks between vendors and clients is another barrier to the accountability of AI systems. Existing governance frameworks view accountability as centralized within the organization deploying an AI system. Such frameworks are lacking explicit means of distributing accountability for different components of system behavior, shared liability for response when vendor AI systems fail, and for bias and performance challenges when crossing organization boundaries [7].

Many governance frameworks do not address continuous monitoring of a vendor's AI systems over their lifecycle but rather point-in-time monitoring at procurement or initial deployment. More mature internal governance of AI systems inside organizations includes active monitoring of model performance over time, distribution drift, correctness, and adherence to specified accuracy and fairness thresholds over the life of the system [8]. However, these systems may not require continuous oversight for vendor AI (as third-party systems are considered static dependencies to be validated at the time of purchase, rather than throughout their operation).

<b>Governance Component</b>	<b>Internal AI Coverage</b>	<b>Vendor AI Coverage</b>	<b>Gap Description</b>
Model lifecycle management	Comprehensive protocols	Point-in-time assessment only	No ongoing monitoring
Data governance structures	Complete visibility	Minimal transparency	Unknown training data composition
Security testing protocols	Comprehensive assessments	Limited vendor disclosure	No adversarial robustness verification
Transparency mechanisms	Full documentation	Limited vendor disclosure	Insufficient behavioral information
Accountability frameworks	Clear internal ownership	Undefined shared responsibility	No collaborative incident response
Continuous oversight	Real-time monitoring	Static dependency treatment	No drift detection capability

Table 2: Critical Omissions in Current AI Governance Frameworks [7, 8]

## 5. Proposed Third-Party AI Governance Framework

### 5.1 Vendor AI Classification System

In third-party AI governance, formal classification schemas categorize vendor AI based on AI risk profiles. Decision criticality assessment describes how vendor AI outputs influence especially consequential decisions for an organization or an individual [9]. Systems that implement AI recommendations without human oversight and those that affect the real world are of the highest importance.

Regulatory impact evaluation assesses vendor AI technology against fairness, transparency, privacy, and safety laws for sectors with meaningful impact. For example, healthcare AI technologies delivering medical devices must be validated and report adverse events. Financial services applications must follow anti-discrimination obligations, while manufacturing AI systems affecting worker safety require compliance with occupational health and safety regulations [10]. Autonomy degree measurement assesses a system's ability to operate without human intervention while data sensitivity categorization estimates the types of data processed and privacy consequences [9].

### 5.2 Shared Accountability Architecture

Beyond vendor management, the framework acknowledges vendors as arm's-length outside parties. The new collaborative AI responsibility also stresses that effective governance relies on active vendor-client partnerships, not on adversarial attitudes that minimize responsibilities [10]. Clear definitions of ownership and liability can delineate which aspects of an AI system's behavior are the vendor's responsibility and which are the client's, including algorithms for processing inputs and generating outputs, security incident response responsibilities, and vulnerability disclosure obligations [9].

### 5.3 Audit-Ready Transparency Standards

Vendor transparency requirements extend beyond documentation. Effective model governance and compliance require full technical transparency on functional capabilities, limitations, input features, architecture, security testing methodologies, and decision-making across various scenarios [9]. Security transparency requirements should include adversarial robustness testing results, vulnerability assessment reports, penetration testing outcomes, and evidence of secure development practices throughout the model lifecycle.

Proactive change management provisions clarify vendor obligations to notify when model updates, architectural changes or retraining are made to a deployed AI system that may affect its behavior or

security posture prior to deploying the changes [10], providing adequate lead time for impact assessment, security validation testing, and integration testing in the deployed application context.

**5.4 Continuous Oversight Mechanisms**

The shift from point-in-time auditing to continuous control of the operational life cycle of vendor AI requires, at its core, real-time monitoring of drift and anomaly patterns, including monitoring of behavioral drifts that indicate performance degradation, bias and drifts related to technical performance [9]. Continuous security monitoring should detect anomalous prediction patterns indicating adversarial attacks, unusual API access patterns suggesting data exfiltration attempts, and performance degradations that may signal model tampering or poisoning.

<b>Framework Pillar</b>	<b>Classification Dimension</b>	<b>Governance Requirement</b>
Vendor AI Classification	Decision criticality	Automatic vs human-mediated decisions
	Regulatory impact	Healthcare, finance, manufacturing, and employment domains
	Degree of autonomy	Fully autonomous vs human-in-the-loop
	Data sensitivity	Protected health, financial, manufacturing, and biometric data
	Security risk level	Critical infrastructure vs general purpose applications
Shared Accountability	Ownership boundaries	Vendor core models vs client deployment
	Risk allocation	Contractual distribution of AI-specific risks
	Security responsibilities	Incident response and vulnerability management
Transparency Standards	Model behavior documentation	Input features and decision patterns
	Security documentation	Adversarial testing and vulnerability assessments
	Change management	Pre-deployment update notification
Continuous Oversight	Output drift monitoring	Real-time behavioral shift detection
	Security monitoring	Anomaly detection and attack identification
	Post-update validation	Automated reassessment triggering

Table 3: Third-Party AI Governance Framework Components [9, 10]

**6. Regulatory Environment and Industry Impact**

Artificial intelligence regulation has risen and matured quickly, and jurisdictions around the world have moved quickly from high-level principles to explicit laws with far-reaching impacts on AI system implementation and governance. Consequently, organizational accountability has become a central focus of regulatory regimes, which in turn often use regulatory language to rebut arguments that organizations

cannot be held liable for the actions of their AI systems by claiming autonomy or the inscrutability of the algorithms [11]. Regulatory frameworks consistently adopt the stance that all entities deploying AI systems are responsible for their outputs and impacts, regardless of technical complexity or algorithmic opacity.

Another emerging regulatory theme is model explainability and interpretability requirements. Such requirements may restrict model designs in some contexts, owing to the need to provide intelligible explanations of AI decision making processes where they impact natural persons' rights, opportunities, or access to services [12]. The explainability imperative presents other challenges for vendor AI adoption, because the information necessary to provide model explainability (e.g. model architectures, feature importance, and decision thresholds) is not always readily available due to vendors' intellectual property and competitive concerns about sharing model details.

Some law converts framework requirements directly into regulation: agents and other entities subject to them are expected to engage in systematic procedural techniques for identifying AI-related risk, evaluating the seriousness and probability of an AI-related risk, implementing control measures to limit unacceptable AI-related risks, and continuously monitoring the effectiveness of the risk-management procedure [11]. The fourth regulatory issue relates to human oversight and intervention: humans should maintain meaningful control over AI systems rather than algorithms acting independently by making decisions with legal or regulatory consequences without human involvement or an opportunity to intervene as appropriate [12].

As a result of third-party AI governance gaps, regulators are increasingly issuing enforcement actions and penalties on organizations operating AI systems that violate applicable law. Organizations have not successfully defended themselves against regulatory penalties by arguing that the problematic AI behavior arose from the AI system delivered by a vendor as opposed to the organization itself, since outsourcing AI capabilities does not provide immunity from liability [11].

Third-party AI governance can also be an opportunity to differentiate competitively, as organizations realize that investing in vendor AI governance beyond the minimum requirements is a source of competitive advantage in cases where a customer, partner, or regulator particularly values fairness, security, and accountability [12].

<b>Regulatory Focus Area</b>	<b>Requirement Type</b>	<b>Vendor AI Challenge</b>	<b>Consequence of Gap</b>
Organizational accountability	Deployer responsibility	No distinction between internal and external	Enforcement actions regardless of origin
Model explainability	Meaningful decision explanation	Limited vendor disclosure	Cannot fulfill explanation obligations
Risk management frameworks	Continuous assessment processes	Point-in-time vendor evaluation	Undetected emerging risks
Human oversight	Meaningful control maintenance	Opaque vendor behavior	Inability to determine intervention points
Security and robustness	Adversarial resilience testing	No access to security assessments	Unmitigated security vulnerabilities
Enforcement actions	Rejection of vendor-blame	Full deployer accountability	Penalties despite the vendor source

	defense		
Competitive differentiation	Proactive governance investment	Strategic positioning advantage	Trust in sensitive sectors

Table 4: Regulatory Focus Areas and Organizational Consequences [11, 12]

**Conclusion**

Third-party AI models are an existing but under-regulated element of contemporary enterprise systems, which pose significant compliance, security, and liability risks with an increasing intensity of regulatory enforcement. The accountability-control paradox places organizations at legal liability for vendor AI systems, the inner mechanics of which are still unknown to them, and the behavior and security posture of which they cannot validate or explain in any meaningful way. The existing governance frameworks focus on internal model management and pay little attention to vendor AI security assessments, so that externally sourced systems could be used outside the risk registers, security testing procedures, control testing protocols, and audit scopes. The suggested governance framework fills this blind spot by proposing the systematic categorization of vendors' AI according to criticality, security risk level, and regulatory influence, shared accountability structures that declare the existence of specific responsibilities for both vendors and clients including security incident response, transparency provisions adequate to address regulatory scrutiny including security testing documentation, and ongoing monitoring systems that become aware of behavioral drift, security anomalies, and performance degradation. Since regulatory bodies are increasingly discarding organizational shields of problematic AI being created by vendors rather than being created internally, proactive control of third-party AI security and governance becomes not only a compliance necessity but a competitive advantage as well. Organizations that have developed extensive AI vendor management demonstrate algorithmic responsibility to stakeholders and position themselves to be responsive as regulatory requirements continue to change based on jurisdictions and localities. As AI systems proliferate across healthcare, financial services, manufacturing, and other critical sectors, the imperative for robust third-party AI governance that addresses both functional and security dimensions will only intensify.

**References**

[1] European Parliament, "EU AI Act: First regulation on artificial intelligence," 2023. [Online]. Available: <https://www.europarl.europa.eu/topics/en/article/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence>

[2] National Institute of Standards and Technology, "Artificial Intelligence Risk Management Framework (AI RMF 1.0)," Jan. 2023. [Online]. Available: <https://www.nist.gov/itl/ai-risk-management-framework>

[3] McKinsey & Company, "The state of AI in 2023: Generative AI's breakout year," 2023. [Online]. Available: [https://www.mckinsey.com/~media/mckinsey/business%20functions/quantumblack/our%20insights/the%20state%20of%20ai%20in%202023%20generative%20ais%20breakout%20year/the-state-of-ai-in-2023-generative-ais-breakout-year\\_vf.pdf](https://www.mckinsey.com/~media/mckinsey/business%20functions/quantumblack/our%20insights/the%20state%20of%20ai%20in%202023%20generative%20ais%20breakout%20year/the-state-of-ai-in-2023-generative-ais-breakout-year_vf.pdf)

[4] Organization for Economic Cooperation and Development, "OECD AI Principles Overview." [Online]. Available: <https://oecd.ai/en/ai-principles>

- [5] Andrew Smith, "Using Artificial Intelligence and Algorithms," Federal Trade Commission, 2020. [Online]. Available: <https://privacysecurityacademy.com/wp-content/uploads/2021/01/Using-Artificial-Intelligence-and-Algorithms--Federal-Trade-Commission.pdf>
- [6] AlgorithmWatch, "Automating Society Report 2020." [Online]. Available: <https://automatingsociety.algorithmwatch.org/>
- [7] The National Institute of Standards and Technology published the "NIST AI RMF Playbook." [Online]. Available: <https://www.nist.gov/itl/ai-risk-management-framework/nist-ai-rmf-playbook>
- [8] Miguel Angel Perez Alvarez et al., "Ethically Aligned Design: A Vision for Prioritizing Human Wellbeing with Artificial Intelligence and Autonomous Systems," ResearchGate, 2017. [Online]. Available: [https://www.researchgate.net/publication/378975517\\_ETHICALLY\\_ALIGNED\\_DESIGN\\_A\\_Vision\\_for\\_Prioritizing\\_Human\\_Wellbeing\\_with\\_Artificial\\_Intelligence\\_and\\_Autonomous\\_Systems](https://www.researchgate.net/publication/378975517_ETHICALLY_ALIGNED_DESIGN_A_Vision_for_Prioritizing_Human_Wellbeing_with_Artificial_Intelligence_and_Autonomous_Systems)
- [9] European Commission, "Regulation Of The European Parliament And Of The Council," 2021. [Online]. Available: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex%3A52021PC0206>
- [10] Incident Database, "AI Incident Database," 2023. [Online]. Available: <https://incidentdatabase.ai>
- [11] European Union Agency for Fundamental Rights, "Getting the Future Right: Artificial Intelligence and Fundamental Rights," Dec. 2020. [Online]. Available: <https://fra.europa.eu/en/publication/2020/artificial-intelligence-and-fundamental-rights>
- [12] Information Commissioner's Office, "Guidance on AI and Data Protection," 2023. [Online]. Available: <https://ico.org.uk/for-organisations/guide-to-data-protection/key-dp-themes/guidance-on-ai-and-data-protection/>