Responsible AI in the Global Context: Maturity Model and Survey

Anka Reuel
Stanford University
Stanford, USA
anka@cs.stanford.edu

Patrick Connolly
Accenture
Dublin, Ireland
patrick.connolly@accenture.com

Kiana Jafari Meimandi
Stanford University
Stanford, USA
kjafari@stanford.edu

Shekhar Tewari
Accenture
Bangalore, India
shekhar.tewari@accenture.com

Jakub Wiatrak
Accenture
Warsaw, Poland
jakub.wiatrak@accenture.com

Dikshita Venkatesh
Accenture
Bengaluru, India
dikshita.venkatesh@accenture.com

Mykel Kochenderfer
Stanford University
Stanford, USA
mykel@stanford.edu

Abstract

Responsible AI (RAI) has emerged as a major focus across industry, policymaking, and academia, aiming to mitigate the risks and maximize the benefits of AI, both on an organizational and societal level. This study explores the global state of RAI through one of the most extensive surveys to date on the topic, surveying 1000 organizations across 20 industries and 19 geographical regions. We define a conceptual RAI maturity model for organizations to map how well they implement organizational and operational RAI measures. Based on this model, the survey assesses the adoption of system-level measures to mitigate identified risks related to, for example, discrimination, reliability, or privacy, and also covers key organizational processes pertaining to governance, risk management, and monitoring and control. The study highlights the expanding AI risk landscape, emphasizing the need for comprehensive risk mitigation strategies. The findings also reveal significant strides towards RAI maturity, but we also identify gaps in RAI implementation that could lead to increased (public) risks from AI systems. This research offers a structured approach to assess and improve RAI practices globally and underscores the critical need for bridging the gap between RAI planning and execution to ensure AI advancement aligns with human welfare and societal benefits.

CCS Concepts

• **Computing methodologies** → **Artificial intelligence**; • **General and reference** → *Surveys and overviews*; • **Security and privacy** → *Human and societal aspects of security and privacy*.

FAccT '25, Athens, Greece

ACM ISBN 979-8-4007-1482-5/25/06
https://doi.org/10.1145/3715275.3732165

ACM Reference Format:
Anka Reuel, Patrick Connolly, Kiana Jafari Meimandi, Shekhar Tewari, Jakub Wiatrak, Dikshita Venkatesh, and Mykel Kochenderfer. 2025. Responsible AI in the Global Context: Maturity Model and Survey. In *The 2025 ACM Conference on Fairness, Accountability, and Transparency (FAccT '25), June 23–26, 2025, Athens, Greece.* ACM, New York, NY, USA, 37 pages. https://doi.org/10.1145/3715275.3732165

1 Introduction

Responsible AI (RAI) has become a central theme in discussions about AI to mitigate risks and harness the benefits of the technology. Recent surveys highlight the growing importance of RAI for organizations worldwide. For instance, a recent report reveals that 52% of companies engage in some level of RAI, albeit with limited scale and scope in most cases [23]. The need for understanding RAI adoption on a company level is increasingly important: RAI measures help identify potential weaknesses that can lead to AI-related harms such as bias, discrimination, and privacy breaches. These issues carry significant societal consequences, potentially exacerbating existing inequalities if not mitigated through widespread organizational and system-level adoption of RAI practices. The absence of standardized metrics impedes the assessment of RAI progress and, consequently, the effective management of AI's societal impact. Regulatory pressure, exemplified by the EU AI Act mandating pre-deployment evaluations for certain AI systems, reflects growing societal concerns about the technology's potential negative effects and necessity for organizations to demonstrate RAI adoption. Demonstrating adherence to RAI practices further builds stakeholder trust and contributes to AI's social acceptance. However, despite its importance, a lack of consensus persists on a standardized, scalable approach to evaluate RAI within organizations. Most existing work [37, 40] focuses on system-level aspects and overlooks organizational considerations, presenting an incomplete picture of RAI adoption, given that organizational processes and plans are a key aspect to ensuring consistency in RAI implementation across different AI projects and teams. This gap in the literature limits our understanding of how RAI principles and frameworks are operationalized and hinders the development

of comprehensive strategies for improving RAI practices. Furthermore, the lack of attention to organizational factors may lead to a disconnect between high-level RAI principles and their practical application in day-to-day operations, potentially undermining the effectiveness of RAI initiatives.

This work addresses these challenges by developing a conceptual RAI maturity model for organizational and system-level RAI maturity. Based on this model, we measure current levels of RAI maturity in a first-of-its-kind global survey with 1000 organizations across 20 industries and 19 geographical regions to track progress and identify areas for improvement in RAI adoption. As part of the survey, we further capture aspects like technical risk mitigation measures, risk perceptions, implemented AI governance structures, and barriers to generative AI adoption to gauge current perceptions regarding the responsible development, deployment, and use of AI, shedding light on how these perceptions might influence RAI maturity. Our aim extends beyond mere quantification of adoption rates. We sought to explore the correlations between RAI adoption, risk exposure, regulatory exposure, and organizations' positions within the supply chain in relation to their RAI activities.

The rest of this work is structured as follows: We outline the background to our work in Sec. 2, our RAI maturity model in Sec. 3, and our survey methodology in Sec. 4. We continue to present the survey results in Sec. 5, before discussing their context and implications in Sec. 6 and limitations in Sec. 7.

2 Background

Responsible AI. Even though RAI is considered a critical aspect in the current AI ecosystem, the concept of RAI lacks a universally accepted definition [37]. Many definitions emphasize the practices used to implement RAI rather than providing a conceptual understanding of its essence [37, 40]. In this work, we base our RAI definition on that of the European Commission's High-Level Expert Group definition of AI principles for system-level RAI dimensions [15]. Their work outlines key aspects of RAI across well-defined categories [28], was developed by a group of subject-matter experts, and was pre-step to the newly enacted EU AI Act.

The European Commission's understanding of RAI encompasses a set of seven system-level (or operational) principles aimed at ensuring ethical and trustworthy AI. We adapted them to account for new developments in the field since the publication of the European Commission's work in 2018 and used the set of system-level RAI dimensions listed in Tab. 1 in our research. The European Commission's definition does not include organizational-level aspects in their RAI understanding. In line with the human-centered AI governance model of Shneiderman [35], we emphasize the importance of organizational processes and governance frameworks, as well as aspects such as training, culture, and leadership support to drive RAI within a company. We add the following dimensions to the RAI understanding used in this study:

- *Organizational Governance*, encompassing structures, processes, and roles within an organization to support RAI efforts [29, 30].
- *Leadership and Culture*, including support from C-level executives, responsibility-first culture, and RAI training efforts [29, 30].

This proposed set of combined dimensions serves as a foundation for the development of our RAI maturity model and RAI survey. It is important to note that these dimensions can often overlap and interact. For instance, transparency is crucial for ensuring fairness and system integrity. Similarly, accountability mechanisms support both reliability and transparency, and in part rely on organizational governance measures to be addressed.

RAI Questionnaires. In recent years, there has been a surge in corporate and governmental efforts to assess RAI. Existing RAI questionnaires and frameworks vary widely, developed by corporations [5, 31], governments [25, 28], and other organizations [1, 33]. These questionnaires differ in type (e.g., context-specific assessment, survey, certification), number of questions, high-level RAI dimensions covered, and methodology. Despite this growing interest, significant gaps exist in current RAI questionnaires and models. Many lack transparent methodologies, fail to define RAI terms or the dimensions they cover, and offer unclear scoring and interpretation. Furthermore, these questionnaires often fail to assess both system-level, i.e., operational, *and* organizational aspects of RAI.

To address the gaps, we specify our definition of the components of RAI, along with the survey structure and maturity logic, and scoring. We further ensure that both system-level mitigation measures and organizational RAI practices are captured by our survey to present the most comprehensive, public RAI survey to date. We further recognize the need for global comparability and have adopted a questionnaire-based approach, which enabled us to scale the survey and gather data from over 1000 organizations worldwide.

3 RAI Maturity Model

As part of the development of our RAI survey, we recognized the lack of and need for an in-depth RAI maturity model that captures how well organizations adopt organizational and system-level, i.e., operational RAI measures and mitigate risks associated with the development, deployment, and use of AI. Maturity models are a well-established tool across fields that provide a straightforward yet powerful way for organizations to assess the quality of their processes[38], with models being developed for software management [18, 27], project management [10], and business processes [13, 16, 20]. For AI, there have been a number of maturity models proposed [4, 11, 14, 34]. However, while there has been interest in a maturity model with a focus on responsible AI [36], existing efforts have been incomplete. Some are limited to system-level RAI maturity models, without accounting for organizational RAI measures [17]. Other efforts have only reviewed literature on (R)AI maturity models without developing their own model based on the identified gaps [2, 3]. Previous studies [2, 21, 24] emphasize the importance of both technical and organizational aspects in RAI maturity, but they do not develop a corresponding comprehensive maturity model themselves. A detailed maturity model is necessary to guide organizations in adopting AI responsibly. This work aims to fill this gap. In addition, we are the first to apply our model in practice, surveying 1000 organizations globally based on our maturity model.

The model comprises two main dimensions: organizational and operational RAI maturity, each with distinct subcomponents and

Component	Description	Question
Reliability	Necessitates that performance is reliable across all relevant (sub-) groups and that AI systems are resilient to attacks, and equipped with safety measures and fallback plans.	Q32
Privacy & Data Governance	Necessitates respect for privacy, data quality, and cybersecurity, including appropriate access controls.	Q29 & Q35
Human Interaction	Protection of fundamental human rights, limitation of misuse, and measures to prevent overreliance on AI systems	Q28
Transparency	AI system that is explainable to the technical user and interpretable to the end-user; also includes release of details about the training data, architecture design, and other relevant system information.	Q33
Societal & Environmental Wellbeing	Promotion of sustainability, positive social contributions, and democratic values.	Q31
Diversity, Non-discrimination & Fairness	AI systems must be developed in a way to actively avoid unfair bias, ensuring accessibility and inclusivity for diverse users through universal design principles and active stakeholder participation in the design process.	Q30
Accountability	Mandating auditability of systems, proactive measures to minimize negative impacts, transparent trade-off discussions, along with mechanisms for recourse and remediation in the event of any adverse effects.	Q34

Table 1: RAI system-level/operational dimensions, based on European Commission [12]. Specific RAI measures for each component are listed in App. B under the question referenced here.

corresponding maturity levels. We iteratively refined both the main and subcomponents and the maturity levels for both based on a synthesis of the literature we outlined above and expert interviews with stakeholders in academia, governance, and industry. The level namings are consistent with standard maturity model naming conventions, as used for example in [14] or [4]. The model presented in Tab. 2, along with the maturity levels defined for the subcomponents (see App. A), can serve as a tool for organizations to identify areas for improvement in adopting RAI practices. It further served as the foundation for our survey design to define the components covered in the survey, as well as the answer options for each question.

3.1 Organizational RAI Maturity

Organizational RAI maturity measures the sophistication and effectiveness of organizational processes, frameworks, and culture in ensuring its responsible development, deployment, and use of AI. Based on previous work (see Tab. 3 for references), we subsume eight subcomponents in organizational RAI maturity, summarized in Tab. 3. Survey questions Q14, Q17, Q20–Q24, and Q33 capture these components in our survey (see App. B). Components and corresponding maturity levels were included based on a literature review for each dimension as well as expert interviews. Scoring details are discussed in Sec. 4.3.

3.2 Operational RAI Maturity

Operational RAI maturity assesses the comprehensiveness of an organization's implementation of system-level mitigation measures for identified AI adoption risks. Operational RAI maturity encompasses all system-level aspects of RAI discussed in Tab. 1. In our survey, we first presented respondents with a list of potential risks (detailed in App. B), including privacy, diversity, reliability, transparency, security, human interaction, environmental impact, and accountability concerns. Respondents selected relevant risks and indicated their level of implementation for pre-defined mitigation measures for these identified risks, such as red teaming, mitigations for adversarial attacks, bias mitigation, or uncertainty quantification (questions Q22b, Q25–Q31, see App. B for all mitigation measures). Options for 'Other' with a free-form text field to add not-listed mitigations measures, and 'None' were also provided. Scoring details are discussed in Sec. 4.3. Mitigation measures were included based on a literature review for each dimension and expert interviews. Scoring details are discussed in Sec. 4.3.

4 Methodology

4.1 Questionnaire Design

In this survey RAI was defined as "the practice of designing, developing, and deploying AI with good intention to empower employees and businesses, and fairly and positively impact customers and society" for the participants.The survey included 39 questions covering system- and organizational-level dimensions of RAI. Ten were qualifier questions about job function, company location, global revenue, visibility into RAI decision-making, industry, and general AI adoption strategy. Respondents had to indicate whether their organization was developing AI models, modifying third-party models, or using third-party models without modification. They also specified if these models were intended for internal use only or for resale. The remaining 29 questions addressed governance, risk management, system-level RAI mitigation measures, talent availability, lawfulness and compliance, and generative AI. The questionnaire was adapted based on responses to the qualifier questions (see App. B). We also considered interviews as a mode for data collection but decided to use a survey to collect data from a broader sample for global representation.

To ensure ethical data collection and informed consent, all participants were provided with clear information about the survey's purpose, scope, and confidentiality measures before participation.

Level	Score	Organizational Maturity	Score	Operational Maturity
Level 1: Initial	[0 , 12.5]	The organization has limited awareness and no organizational plans, processes, or frameworks in place to ensure a responsible AI adoption.	[0 , 12.5]	The organization does not mitigate identified risks on a system level.
Level 2: Assessing	[12.5 , 37.5]	The organization is aware of the necessity for organizational measures to ensure a responsible AI adoption and is assessing governance options.	[12.5 , 37.5]	Awareness of risks may be present, but the organization has only limited or no formal mitigation measures in place.
Level 3: Determined	[37.5 , 62.5]	The organization demonstrates foundational governance capabilities to support the responsible development, deployment, and use of AI.	[37.5 , 62.5]	A few risk mitigation measures are being fully operationalized, but the majority is only implemented ad-hoc or in early roll-out stages. There is a growing awareness of the need for more systematic approaches.
Level 4: Managed	[62.5 , 87.5]	The organization has established comprehensive organizational RAI measures and is actively ensuring enterprise-wide adoption, demonstrating a mature and effective approach to internal RAI governance.	[62.5 , 87.5]	A wide range of risk mitigation measures are fully operationalized across all relevant AI systems in the organization.
Level 5: Optimized	[87.5 , 100]	The organization demonstrates an established, future-oriented approach towards organizational RAI, ensuring a sustainable and responsible approach to organizational RAI.	[87.5 , 100]	Comprehensive, state-of-the-art risk mitigation strategies are fully operationalized. The organization continuously monitors and evaluates risks, proactively adapting its practices as needed to mitigate new risks.

Table 2: Overall organizational and operational maturity model and corresponding score brackets used for our global RAI survey. Individual maturity models per subcomponent can be found in App. A.

Participants were notified that the survey would take approximately 25-30 minutes to complete and were explicitly informed about data confidentiality - including that their identities and responses would be kept strictly confidential and that all responses would be aggregated without attribution to specific individuals or companies. Participation was entirely voluntary, and respondents could withdraw at any time. This consent process was integrated into the survey platform's double opt-in recruitment process, requiring explicit acknowledgment before proceeding with the survey.

4.2 Data Collection

To ensure the relevance of the responses, we included only C-suite respondents who were directly involved in or had visibility into RAI-related strategies and priorities. We targeted companies actively developing, selling, or using AI commercially, focusing on those identifying specific AI risks or anticipating AI regulation within five years. Companies with less than $499 million USD in global annual revenue or where respondents lacked visibility into the RAI decision-making process were excluded. An outside firm was contracted to manage recruitment and data collection, using a double opt-in process and various recruitment channels such as LinkedIn and partnerships. Quality control measures were implemented to ensure data integrity, including screening for low-quality responses, which were identified by unusually short completion times or patterns of uniform answers. The study was designed to ensure global representation by selecting 20 countries across all major geographical regions. We received 1,000 complete responses across 19 industries. To accommodate non-English-speaking respondents, surveys were translated into official languages of the respective countries. Data collection occurred between February and March 2024, with participation being anonymous and voluntary. Informed consent was obtained, and respondents had the option to withdraw at any point. Participants were incentivized through the external firm's incentive system.

4.3 Data Analysis & Scoring

For questions assessing organizational processes and frameworks (Q14, Q17, Q20–Q24, Q33), we formulated the answer options in our

Sub Component	Description	Source
	Organizational Governance	
Governance Operating Model	Structures, policies, and processes that guide AI system development and deployment.	Batool et al. [6], Lu et al. [22]
Risk Identification	Systematic processes to recognize the potential risks and developing proactive measures to detect problems in AI systems before they escalate.	Clarke [9], Lu et al. [22], Qureshi et al. [32]
Risk Assessment & Management Framework	Comprehensive risk management framework and processes for the evaluation of the likelihood and potential impact of the identified risks to prioritize their management.	Clarke [9], Wirtz et al. [39]
Risk Mitigation	Processes to implement strategies and controls to reduce the potential negative impacts of AI systems; Includes frameworks to address potential negative impacts on individuals, the organization, and society.	Clarke [9], Qureshi et al. [32]
Procurement	Processes involved in acquiring AI technologies and services from external vendors; Ensuring procurement decisions align with RAI principles.	Obinna and Kess-Momoh [26]
Monitoring & Control	Ongoing oversight of AI systems to ensure they operate as intended and adhere to the guidelines and mechanisms for detecting and correcting deviations.	Camilleri [8]
Cybersecurity	Protection of AI systems and data from cyberthreats. Implementing robust security measures to safeguard against unauthorized access, data breaches, and other cyberrisks.	Bowen et al. [7], Lannquist[1] et al. [19]
	Leadership and Culture	
Sponsorship	Support and commitment from senior leadership for RAI initiatives and allocating resources, setting priorities, and endorsing RAI practices.	Lu et al. [22]
Training	Education and training of employees and stakeholders on RAI practices.	

Table 3: RAI organizational dimensions and their sub components. Detailed maturity levels for each subcompontent can be found in App. A.

survey so that they corresponded, for the respective subcomponent assessed, to the organizational maturity levels outlined in Tab. 2, with each level scoring 0, 25, 50, 75, or 100 points respectively for each subcomponent (see App. A for all subcomponent-specific maturity levels). The overall organizational RAI maturity score was derived by averaging the scores from these questions, providing a measure of a company's organizational RAI maturity. For questions assessing concrete measures taken within a system-level RAI dimension (Q22b, Q25–Q31), we assigned 100 points if a mitigation measure has been indicated to be fully operationalized and 0 otherwise. We only asked about measures for risks that the respondents previously identified as relevant to their AI adoption. For operational RAI maturity, we averaged the points across all relevant mitigation measures and assigned a maturity level based on the resulting average score (see Tab. 2 for scoring brackets). While some adopted measure may be more important than others, we opted for a simple average, rather than a weighted average, because the (perceived or actual) importance of these measures is often context-dependent and a weighting system would not be objectively justifiable at this level of abstraction. The scores were further aggregated at multiple levels (e.g., region, industry) to enable comparative analyses (see App. C). Independent samples t-tests ($\alpha = 0.05$) were employed to assess statistical significance of group differences. We further ran regression analysis between selected questions (see App. C.0.1).

5 Results

5.1 Maturity

5.1.1 Organizational RAI Measures. Most organizations are at the *Managed* stage of organizational RAI maturity, with only 9% reaching the *Optimized* stage, indicating that fully integrated RAI practices are still uncommon. However, the majority of organizations are at mid-level organizational RAI maturity which can be interpreted as an indication of a broad recognition of RAI while highlighting challenges in advancing beyond this stage (Fig. 1). Maturity varies by region, with North America and Asia leading, particularly Singapore and Japan. Latin America lags, with many organizations at the *Initial* stage and the lowest average organizational RAI maturity score (see App. C). Across the selected 19 industries, *Healthcare*, *Life Sciences*, and the *Communication, Media, & Technology* sectors exhibit high RAI maturity, likely driven by higher regulatory demands and technological innovation within these industries. In contrast, *Natural Resources* and *Utilities* show lower maturity, highlighting a need for improvement (see App. C).

5.1.2 Operational RAI Measures. The operational RAI index results in a right-skewed distribution, with a mean score around 35. A floor

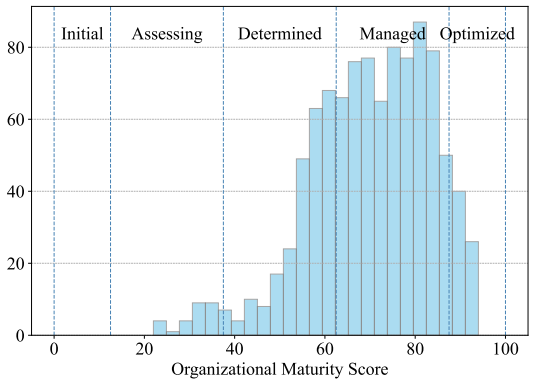

Figure 1: Organizational RAI maturity distribution.

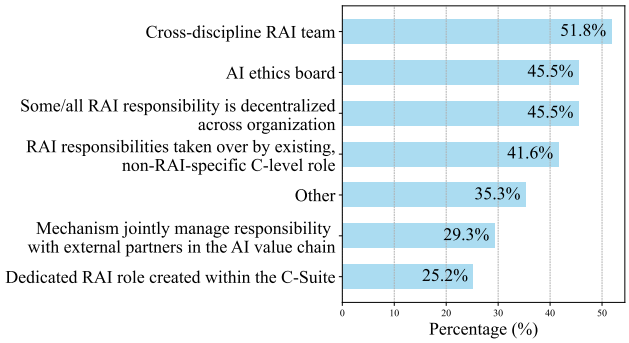

Figure 2: Reported RAI roles and structures within organizations.

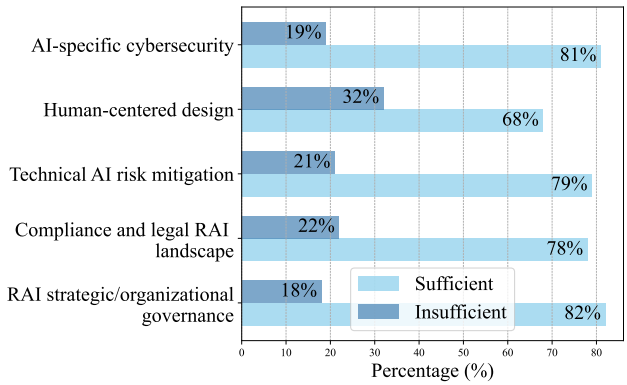

Figure 3: Indicated availability of sufficient talent with specific RAI skills.

effect shows 7% of respondents scored 0, indicating no operationalized RAI measures. The majority of organizations are at a mid-level of operational maturity while only a small fraction of organizations reach the *Optimized* stage, indicating that achieving full operational maturity in RAI is challenging and potentially under-prioritized. Regionally, mean operational maturity scores are similar, suggesting uniform global operational RAI maturity. However, Latin America has the highest proportion of companies with no operationalized RAI measures, indicating a lag in adoption. North America and Europe show the lowest proportions, reflecting more widespread RAI implementation and a higher baseline maturity.

5.2 Individual Questions

5.2.1 Risk Perception. The top risks associated with AI were privacy and data governance (51%), cybersecurity (47%), and reliability (45%) concerns. These highlight the importance of managing sensitive data and ensuring secure, robust AI operations. In contrast, non-discrimination and fairness (29%) and accountability (26%) concerns were seen as less urgent. Organizational/business (12%) and brand risks (26%) were among the least selected, indicating a lower perceived impact on operations and direct concern for reputation. Geographically, organizations vary in the number of risks they identify, with Asian companies selecting the most (4.99), followed by Rest of the World (4.56), Europe (4.32), and North America (4.19), reflecting regional differences in risk perception. A significant majority (88%) of organizations believed the responsibility for mitigating risks specifically from generative AI lies with foundation model developers, not end-users. Additionally, human interaction risks (35%) and environmental issues (30%) were increasingly recognized, reflecting growing concerns about the societal and ecological impacts of AI.

5.2.2 View of RAI. Participants were asked about how they view RAI within their organization (Fig. 9). Out of all respondents, 49.3% indicated that they see RAI as a strategic tool for revenue growth while 46.2% answered that it is a way to improve the performance of their AI models/systems. Furthermore, 43% saw it as a way to improve their brand reputation and trustworthiness, and 42.5% saw it as a necessity to ensure the safety and security of their AI systems. In comparison, 13.1% indicated that it was not critical for their current usage of AI systems, while only 13.4% indicated it is slowing down innovation and time to market. These results highlight that overall, participants saw RAI as a value driver, and only a minority indicated that it would not be relevant or negatively affect their operations.

5.2.3 Roles & Structures. The majority of organizations (51.8%) indicated that they have a cross-discipline RAI team, followed by 45.5% of respondents saying that some or all of their RAI responsibility is decentralized across business units and functions (Fig. 2). 45.5% indicated that they have an AI ethics board (or equivalent), while only 25.2% indicated having a dedicated C-level RAI role.

5.2.4 Sufficient Talent for Key RAI Skills. Companies were further asked to indicate if they have sufficient talent currently available to them for key RAI skills. The results present a relatively uniform picture, with about 20% of companies indicating that they have insufficient people with AI-specific cybersecurity, technical AI risk mitigation, legal & regulatory RAI, and strategic or organizational RAI governance skills. The exception is human-centered design, where 32% of organizations indicated that they do not have enough talent with this skill set to meet current demands.

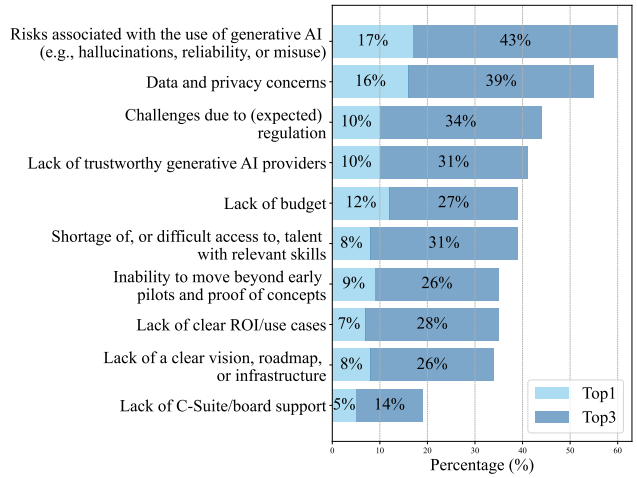

Figure 4: Barriers to the use and development of generative AI.

5.2.5 Barriers to Generative AI Adoption. When it comes to barriers to adopting generative AI, organizations prioritize harms (43%), data and privacy concerns (39%), and regulatory challenges (34%) over traditional hurdles such as talent acquisition (31%), budget constraints (27%), and scalability issues (26%). This shift highlights a growing awareness of the unique challenges posed by generative AI.

6 Discussion

Lack of Responsible AI Maturity. No organization has reached the *Optimized* stage in both organizational and operational dimensions. 9% have achieved *Optimized* organizational maturity, but only 0.8% have reached operational maturity (Fig. 6), indicating a gap between planning and execution of RAI practices. This suggests formal RAI structures and policies exist, but implementation lags. The discrepancy between organizational and operational maturity poses a societal risk. Organizations might appear more prepared to handle AI responsibly than they actually are in practice. This could lead to a false sense of security among stakeholders and potentially inadequate safeguards against AI-related risks. The operational distribution's left skew indicates widespread difficulties implementing RAI practices, reflecting challenges in translating RAI principles and frameworks into operations. A statistically significant negative association exists between the number of relevant risks and operational maturity score ($p < 0.05$, App. C.0.1). There's also a negative correlation between regulatory exposure and both maturity types. This suggests that increased risk and regulatory exposure complicates system-level mitigation efforts, a potential indicator for the complexity of RAI implementation efforts.

Responsibility Uncertainties. The adoption of generative AI is significantly impeded by unresolved risks and perceived regulatory challenges. Yet, companies developing foundational models are seen as responsible for mitigating associated risks, rather than the organizations utilizing these models. This perspective is reflected

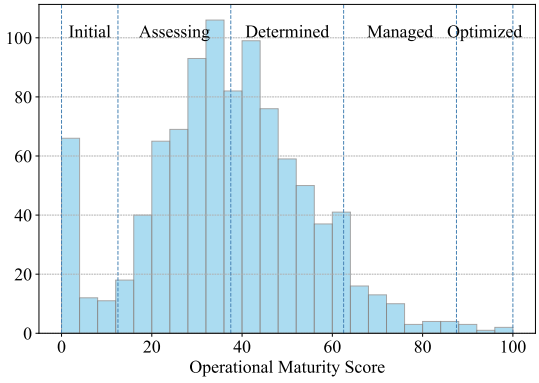

Figure 5: Operational RAI maturity distribution.

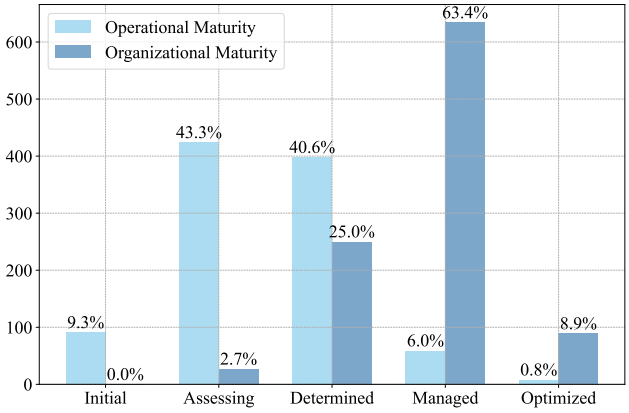

Figure 6: Comparison of organizational and operational maturity levels across all the organizations.

in the operational maturity distribution, where a significant portion of entities cluster in the *Initial* and *Assessing* categories, potentially indicating a hesitancy among organizations to fully adopt mitigation measures due to unclear risk management responsibilities. Existing regulations, such as the EU AI Act and consumer protection laws, should be examined to clarify these responsibilities and ensure alignment with regulatory expectations. This regulatory clarity could potentially shift the entire distribution rightward, as organizations would have clearer guidelines for RAI development and deployment.

Humans at the center of RAI. Companies lack sufficient talent for effective RAI activities, especially in human-centered design, with human interaction being the least developed dimension in operational RAI maturity. 72% of companies further indicated employees and end-users are critical in identifying and mitigating risks like hallucinations, cybersecurity threats, and IP/data breaches. However, ongoing RAI training for employees is limited, leaving them unprepared for evolving risks, potentially increasing public safety risks. The human-centered design skills shortage could result in AI

systems misaligned with human values and ethics, potentially harming vulnerable populations or reinforcing societal biases. This might erode trust in AI systems, slowing adoption and the realization of societal benefits. The gap between the recognized importance of employee involvement and the lack of ongoing training may create challenges for regulators. It could be difficult to enforce RAI practices when the workforce lacks the necessary skills and knowledge to do so.

7 Limitations

The study only included companies with annual global revenue exceeding 499 million USD, potentially excluding smaller but relevant entities. The survey provides a sampled perspective rather than complete, representative data of the population of interest, and responses are inherently subjective, with potential cultural influences affecting uniformity. Surveys tend to be subject to a positive response bias compared to other means of data collection. Design decisions, such as not providing a 'Not Applicable' option for some questions, could have impacted responses, too. The survey methodology is subject to biases, including social-desirability bias, convenience sampling, non-response bias, and self-selection bias, leading to a sample that may not accurately represent the target population. Techniques such as anonymity, pre-testing, indirect questions, response timing, and reverse-coding were employed to minimize these effects. These limitations underscore the need for cautious extrapolation of the findings to the broader population.

8 Conclusion

Despite progress towards RAI maturity, significant gaps in organizational and operational implementation persist, potentially exposing society to unmitigated AI risks by organizations developing, deploying, and using AI. These gaps could lead to inadequate safeguards and AI systems misaligned with broader societal values. The difficulty in translating ethical principles into practice, especially in human-centered design and interaction, coupled with limited employee training on evolving AI risks, presents critical societal challenges. These issues could lead to increased public safety risks and eroded trust in AI systems, potentially slowing AI adoption and its societal benefits. Addressing these challenges requires collaborative efforts from industry, academia, and policymakers to bridge the gap between RAI planning and execution, focusing on human-centered approaches and comprehensive risk mitigation strategies to ensure AI advancement aligns with societal values.

Acknowledgments

We specifically want to highlight the contributions of Accenture Chief Responsible AI Officer Arnab Chakraborty and the Accenture Research team to the data collection and the financial support of this collaboration.

References

[1] AI Global. 2020. Responsible AI Design Assistant. https://oproma.github.io/rai-trustindex/. Accessed: January 10, 2023.

[2] Pouria Akbarighatar. 2022. Maturity and readiness models for Responsible Artificial Intelligence (RAI): a systematic literature review.

[3] Pouria Akbarighatar, Ilias Pappas, and Polyxeni Vassilakopoulou. 2023. A sociotechnical perspective for responsible AI maturity models: Findings from a mixed-method literature review. *International Journal of Information Management Data Insights* 3, 2 (2023), 100193.

[4] Sulaiman Alsheibani, Yen Cheung, and Chris H Messom. 2019. Towards An Artificial Intelligence Maturity Model: From Science Fiction To Business Facts.. In *PACIS*. 46.

[5] appliedAIm. 2023. AI Maturity Assessment Tool. https://www.appliedai.de/maturity-assessment. Accessed: January 11, 2023.

[6] Amna Batool, Didar Zowghi, and Muneera Bano. 2023. Responsible AI governance: a systematic literature review. *arXiv preprint arXiv:2401.10896* (2023).

[7] Gordon Bowen, Janakan Sothinathan, and Richard Bowen. 2024. Technological Governance (Cybersecurity and AI): Role of Digital Governance. In *Cybersecurity and Artificial Intelligence: Transformational Strategies and Disruptive Innovation*. Springer, 143–161.

[8] Mark Anthony Camilleri. 2024. Artificial intelligence governance: Ethical considerations and implications for social responsibility. *Expert systems* 41, 7 (2024), e13406.

[9] Roger Clarke. 2019. Principles and business processes for responsible AI. *Computer Law & Security Review* 35, 4 (2019), 410–422.

[10] J Kent Crawford. 2021. *Project management maturity model*. Auerbach Publications.

[11] Anna Paula Tanajura Ellefsen, Joanna Oleśków-Szłapka, Grzegorz Pawłowski, and Adrianna Toboła. 2019. Striving for excellence in AI implementation: AI maturity model framework and preliminary research results. *LogForum* 15, 3 (2019).

[12] European Commission. 2019. High-level expert group on artificial intelligence. *Ethics guidelines for trustworthy AI* 6 (2019).

[13] David M Fisher et al. 2004. The business process maturity model: a practical approach for identifying opportunities for optimization. *Business Process Trends* 9, 4 (2004), 11–15.

[14] Philipp Fukas, Jonas Rebstadt, Florian Remark, and Oliver Thomas. 2021. Developing an Artificial Intelligence Maturity Model for Auditing.. In *ECIS*.

[15] High-Level Expert Group on Artificial Intelligence. 2018. A definition of AI: Main capabilities and scientific disciplines. https://ec.europa.eu/.

[16] David A Hillson. 1997. Towards a risk maturity model. *The international journal of project & Business risk management* 1, 1 (1997), 35–45.

[17] Marianna Jantunen, Erika Halme, Ville Vakkuri, Kai-Kristian Kemell, Rebekah Rousi, Tommi Mikkonen, Anh Nguyen Duc, and Pekka Abrahamsson. 2021. Building a maturity model for developing ethically aligned AI systems. *IRIS (Information Systems Research in Scandinavia)* (2021).

[18] Meenu Mary John, Helena Holmström Olsson, and Jan Bosch. 2021. Towards mlops: A framework and maturity model. In *2021 47th Euromicro Conference on Software Engineering and Advanced Applications (SEAA)*. IEEE, 1–8.

[19] Yolanda Lannquist[1], Jia Yuan Loke[1], Nicolas Miailhe[1], Cyrus Hodes[1], and Roman V Yampolskiy. 2020. The intersection and governance of artificial intelligence and cybersecurity. *The Future Society* (2020).

[20] Jihyun Lee, Danhyung Lee, and Sungwon Kang. 2007. An overview of the business process maturity model (BPMM). In *Asia-Pacific Web Conference*. Springer, 384–395.

[21] Ulrich Lichtenthaler. 2020. Five maturity levels of managing AI: From isolated ignorance to integrated intelligence. *Journal of Innovation Management* 8, 1 (2020), 39–50.

[22] Qinghua Lu, Liming Zhu, Xiwei Xu, Jon Whittle, Didar Zowghi, and Aurelie Jacquet. 2024. Responsible AI pattern catalogue: A collection of best practices for AI governance and engineering. *Comput. Surveys* 56, 7 (2024), 1–35.

[23] MIT Sloan Management Review and the Boston Consulting Group. 2022. New Report Documents Business Benefits of Responsible AI. https://mitsloan.mit.edu/ideas-made-to-matter/new-report-documents-business-benefits-responsible-ai. Accessed: July 2, 2024.

[24] Subhodeep Mukherjee and Venkataiah Chittipaka. 2022. Analysing the adoption of intelligent agent technology in food supply chain management: an empirical evidence. *FIIB Business Review* 11, 4 (2022), 438–454.

[25] National Institute of Standards & Technology. 2023. Artificial Intelligence Risk Management Framework (AI RMF 1.0).

[26] Amaka Justina Obinna and Azeez Jason Kess-Momoh. 2024. Developing a conceptual technical framework for ethical AI in procurement with emphasis on legal oversight. *GSC Advanced Research and Reviews* 19, 1 (2024), 146–160.

[27] Mark C Paulk. 1994. A Comparison of ISO 9001 and th Capability Maturity Model for Software. *Technical Report CMU/SEI-94-TR-12 ESC-TR-94-12* (1994).

[28] AP Pekka, W Bauer, U Bergmann, M Bieliková, C Bonefeld-Dahl, Y Bonnet, L Bouarfa, et al. 2018. The European Commission's High-Level Expert Group on Artificial Intelligence. Ethics guidelines for trustworthy AI.

[29] PricewaterhouseCoopers International Limited. 2019. *A practical guide to Responsible Artificial Intelligence (AI)*. Technical Report. PricewaterhouseCoopers International Limited.

[30] PricewaterhouseCoopers International Limited. 2021. *Responsible AI - Maturing from Theory to Practice*. Technical Report. PricewaterhouseCoopers International Limited.

[31] PricewaterhouseCoopers International Limited. 2023. *Responsible AI Diagnostics Tool.* Technical Report. PricewaterhouseCoopers International Limited.
[32] Naila Iqbal Qureshi, Apeksha Garg, Preeti Singh, and Niklas Retzlaff. 2024. AI and Corporate Risk Management: Identifying and Mitigating Technological and Ethical Risks. In *2024 International Conference on Knowledge Engineering and Communication Systems (ICKECS)*, Vol. 1. IEEE, 1–5.
[33] Responsible Artificial Intelligence Institute. 2021. A Certification for Responsible AI. https://assets.ctfassets.net. Accessed: February 1, 2023.
[34] Raghad Baker Sadiq, Nurhizam Safie, Abdul Hadi Abd Rahman, and Shidrokh Goudarzi. 2021. Artificial intelligence maturity model: a systematic literature review. *PeerJ Computer Science* 7 (2021), e661.
[35] Ben Shneiderman. 2020. Bridging the gap between ethics and practice: guidelines for reliable, safe, and trustworthy human-centered AI systems. *ACM Transactions on Interactive Intelligent Systems (TiiS)* 10, 4 (2020), 1–31.
[36] Ville Vakkuri, Marianna Jantunen, Erika Halme, Kai-Kristian Kemell, Anh Nguyen-Duc, Tommi Mikkonen, and Pekka Abrahamsson. 2021. Time for AI (ethics) maturity model is now. *arXiv preprint arXiv:2101.12701* (2021).
[37] Qiaosi Wang, Michael Madaio, Shaun Kane, Shivani Kapania, Michael Terry, and Lauren Wilcox. 2023. Designing responsible ai: Adaptations of ux practice to meet responsible ai challenges. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*. 1–16.
[38] Roy Wendler. 2012. The maturity of maturity model research: A systematic mapping study. *Information and software technology* 54, 12 (2012), 1317–1339.
[39] Bernd W Wirtz, Jan C Weyerer, and Ines Kehl. 2022. Governance of artificial intelligence: A risk and guideline-based integrative framework. *Government Information Quarterly* 39, 4 (2022), 101685.
[40] Boming Xia, Qinghua Lu, Harsha Perera, Liming Zhu, Zhenchang Xing, Yue Liu, and Jon Whittle. 2023. Towards Concrete and Connected AI Risk Assessment (C 2 AIRA): A Systematic Mapping Study. *2023 IEEE/ACM 2nd International Conference on AI Engineering–Software Engineering for AI (CAIN)* (2023), 104–116.

A Maturity Levels for Subcomponents

Tab. 4 and Tab. 5 below describe the maturity levels that we developed and that were used as part of the survey.

B Survey Questions

This section contains the questions that respondents were asked as part of the survey. If the answer options are enumerated, respondents could only select one answer option. The enumeration is missing in case of questions where respondents could select multiple answers. Information in brackets [ ] describes the survey logic, e.g., what answers led to survey terminations (e.g., due to an exclusion based on the qualifier questions) or which questions where only asked if a precondition was fulfilled. For some question, a scale was indicated and respondents had to select, for each subquestion listed, the answer selecting from the defined scale.

q1. Which of the following most closely matches your current job function?

- Board Member
- Chief Executive Officer (including business units & geographic/market CEOs)
- Chief Compliance Officer (or equivalent)
- Chief Information Security/Cyber Security Officer (or equivalent)
- Chief Risk Officer
- Chief Digital Officer
- Chief Analytics/AI Officer/Chief Data and Analytics Officer/Chief Data Officer/ or equivalent
- Chief Information Officer/Chief Technology Officer
- SVP/VP/Director of AI/ML Engineering or equivalent (reporting to any of the above titles)
- Chief Data Scientist
- Chief Legal Counsel (or equivalent)
- Other [Terminate]

q2. In which country is your organization headquartered?

- United Arab Emirates
- Argentina
- Australia
- Brazil
- Canada
- China
- Germany
- Denmark
- Spain
- Finland
- France
- United Kingdom
- India
- Italy
- Japan
- Mexico
- Norway
- Saudi Arabia
- Sweden
- Singapore
- United States
- South Africa
- Other [Terminate]

Sub Component	Level 1	Level 2	Level 3	Level 4	Level 5
Governance Operating Model	None/Use of existing governance operating model with no/limited RAI focus	RAI-specific governance framework designed, following gap analysis of existing governance model against AI risks	RAI-specific framework translated into organization-wide operating model with RAI-specific roles and responsibilities	RAI-specific governance operating model fully operationalized across the organization, with clear RAI roles and responsibilities, processes, reporting, tracking, KPIs and incentives for front- line personnel	Fully-operationalized RAI-specific governance model that is extended beyond the organization to incorporate partners in the supply chain/ecosystem to ensure overarching responsibility with respect to their AI models/systems
Risk Management	No risk management framework in place.	Use of the existing, non-AI-specific risk management framework.	Performed a gap analysis of the existing risk management processes completed against all major AI risks, but no defined RAI risk management strategy and/or framework yet.	Defined or rolling out an RAI risk management strategy and an organization-wide RAI risk management framework.	Fully operationalized RAI risk management framework across the organization and dedicated resources to ensure the approach stays up to date in light of advances in AI, new RAI risk management best practices and policy developments.
Risk Identification	No approach to identifying AI risks beyond performance metrics (e.g., accuracy) and no awareness of the risks.	Awareness of the applicable risk(s) but no structured approach to identifying them.	Defined or rolling out a structured approach to risk identification, based on an analysis of the intended purpose of the AI model/system, foreseeable/-known forms of misuses and potentially affected stakeholders.	Fully operationalized structured approach to risk identification.	Fully operationalized structured approach to risk identification, with dedicated resources in place to proactively anticipate new risks.
Risk Mitigation	No processes in place to mitigate identified RAI risks.	Ad-hoc mitigation techniques, typically led by data scientists/RAI practitioners at project level.	Structured processes to integrate some mitigation techniques into existing AI development processes.	Structured processes to integrate a comprehensive set of mitigation techniques into existing AI development processes with clear stage gates, testing/evaluation procedures and multi-disciplinary reviews.	Structured processes to integrate a comprehensive set of mitigation processes and techniques into existing AI development processes with clear stage gates, testing/evaluation procedures, multi-disciplinary reviews, and proactive research on new advances in risk mitigation.

Table 4: Maturity levels for organizational RAI subcomponents.

Sub Component	Level 1	Level 2	Level 3	Level 4	Level 5
Monitoring & Control	No specific monitoring and control process in place	Some general (non-RAI specific) monitoring by MLOps and/or similar teams within the organization	Centralized plan for RAI monitoring and control defined/rolling out to monitor and identify risks in production, with defined cross-domain responsibilities, roles, and processes	RAI monitoring and control system fully operationalized to monitor and identify risks in production, with defined cross-domain responsibilities, roles, and processes	Monitoring and control systems fully operationalized and optimized to proactively and continuously identify new RAI risks
Cybersecurity	No awareness of any AI-related cybersecurity risks and no related measures in place.	Awareness of AI-related cybersecurity risks but no related cybersecurity measures in place.	Use of the existing cybersecurity policies in an ad-hoc manner to address AI-related incidents.	Clearly outlined AI-specific cybersecurity incident response plan for containing, investigating, and rectifying AI-related events.	Clear AI-specific cybersecurity incident response plan with vulnerability management measures for model audits, adversarial testing, etc., within the AI risk-management framework.
Sponsorship	No active/explicit involvement	RAI acknowledged, but not yet part of broader organization strategy	Advocates on RAI, and is involved in some RAI efforts, e.g., initiating RAI principles	Takes full responsibility on RAI and leads efforts, e.g., setting up RAI teams, overseeing implementation of organization-wide RAI policies and specifying RAI-related C-level KPIs	Takes full responsibility on RAI and leads efforts, while also working with ecosystem of peers/experts/regulators to evolve RAI initiatives
Training	No official RAI training currently available	Ad-hoc RAI training in siloes/by project	Guidance and high-level RAI training provided for those holding RAI responsibilities, including end users, business owners, data scientists, etc.	RAI training is part of all AI-related employee training. One-time role-specific RAI training provided to all employees with RAI responsibilities, including end users, business owners, data scientists, etc.	RAI training is part of all AI-related employee training. Ongoing role-specific RAI training provided to all employees with RAI responsibilities including end users, business owners, data scientists, etc.

Table 5: Maturity levels for organizational RAI subcomponents (continued).

q3. In what countries does your organization currently produce (i.e., design, develop, implement), use or sell AI? Select all that apply.

- United Arab Emirates
- Argentina
- Australia

- Brazil
- Canada
- China
- Germany
- Denmark
- Spain
- Finland
- France
- United Kingdom
- India
- Italy
- Japan
- Mexico
- Norway
- Saudi Arabia
- Sweden
- Singapore
- United States
- South Africa
- Other [Terminate IF ONLY]
- None; we don't produce, use, or sell AI [Exclusive, Terminate]

q4. Which industry does your organization mainly operate in?

- Aerospace & Defense
- Automotive
- Banking
- Capital Markets
- Chemicals
- Telecommunications, Media & Entertainment
- Consumer Goods & Services
- Energy
- Healthcare
- High Tech
- Industrial Equipment
- Insurance
- Life Sciences
- Natural Resources
- Public Services
- Retail
- Software & Platforms
- Travel & Transport
- Utilities
- Other [Terminate IF ONLY]

q5. What was the total annual GLOBAL revenue of your organization in the latest financial year (including all regions and business units)?

- Less than $1 million [Terminate]
- $1 to $49 million [Terminate]
- $50 to $99 million [Terminate]
- $100 to $499 million [Terminate]
- $500 to $999 million
- $1 to $4.9 billion
- $5 to $9.9 billion
- $10 to $19.9 billion
- $20 to $49.9 billion
- $50 billion or more

q6. Which of the following statements best describes your role regarding the development and implementation of, and/or compliance with responsible AI (RAI) related strategies and priorities within your organization?

- We currently do not have a specific RAI strategy or plan in place
- Informed only after the strategy is finalized with responsibilities to execute on the plan
- Acting in the capacity of an advisor to the C-suite and the board in these aspects, but do not have any influence on the final strategy
- Fully involved and working in collaboration with the C-suite and the board, business, and functions in setting the vision, performance criteria, KPIs, etc., but do not share direct ownership and responsibility
- Fully involved and working in collaboration with the C-suite and the board, business, and functions in setting the vision, performance criteria, KPIs, etc., and share direct ownership and responsibility
- Not involved or have any visibility to the RAI decision-making process [Terminate]

q7. How many AI models/systems are you currently using/developing or planning to use/develop in the next two years in your organization?

- 0
- 1
- 2 to 5
- 6 to 10
- 11 to 20
- 21 or more
- Currently using/developing
- Using/developing in next 2 years [Terminate IF = 1]

q8. Which of the following statements best describes your current AI adoption and implementation strategy?

- We do not use AI and an AI adoption strategy is irrelevant to us. [Terminate]
- Organization has an ad-hoc AI strategy with nascent data practices, model development/procurement processes, etc.
- Organization has defined/is rolling out an AI strategy, covering key areas like governance, roles, data, model development/procurement, tooling, MLOps, talent, partnerships, etc.
- Organization-wide AI strategy is fully operationalized, with a sufficiently responsible data strategy, centrally managed tools, and basic AI procurement and talent strategy.
- Organization has processes in place to continuously evolve and optimize all components of its operationalized AI strategy.

q9. Which of the following risks are relevant to your current and future AI models/systems? Select all that apply.

- Diversity & Non-discrimination risks (e.g., fairness concerns, toxicity, discrimination and stereotype reproduction)
- Privacy & Data Governance risks (e.g., data leakage, unauthorized usage of data, etc.)

- Reliability risks (e.g., output errors, hallucinations, model failure)
- Security risks (e.g., cybersecurity incidents)
- Human Interaction risks (e.g., misuse by users for the generation of deepfakes or misinformation, overreliance of users/employee on AI models/systems, or physical/mental harm due to model/system usage)
- Transparency risks (e.g., inexplicable model decisions)
- Societal risks (e.g., threats due to (semi-) autonomous decisions, threats to political stability, national security concerns)
- Environmental risks (e.g., high carbon footprint of model training, inference, hardware)
- Client/Customer risks (e.g., loss of trust, market share, customer satisfaction)
- Brand/Reputational risks (e.g., damage caused to brand by AI-related incident)
- Accountability risks (e.g., IP/copyright violations)
- Compliance and Lawfulness risks (e.g., incompliance with AI regulations or related laws)
- Organizational/Business risks (e.g., lack of AI ROI, AI-related financial loss)
- Other (please specify)
- None [Exclusive]

q10. Do you expect your organization to be subject to any specific regulation of AI/legal liabilities over the next 5 years, based on your current AI strategy, adoption, and geographical locations?

- Yes, we are already subject to such regulation of AI/legislations
- Yes, within 3 years
- Yes, within 5 years
- No [Terminate IF q9 = 15]
- Unsure

q11. Please list the (types of) regulation of AI/legal liabilities that you feel apply/will apply to your organization (e.g., EU AI Act, New York AI Bias Law, EU Liability Directive, future national-level AI regulation, future AI-cybersecurity laws). *[Free text, Optional]*

q12. Please list the (types of) regulation of AI/legal liabilities that you feel could apply to your organization (e.g., EU AI Act, New York AI Bias Law, EU Liability Directive, future national-level AI regulation, future AI-cybersecurity laws). *[Free text, Optional]*

q13. Looking first at Generative AI, what is your organization's current and future (i.e., within two years) adoption strategy?

Current. Select all that apply.

- Build and Use: Develop your own (proprietary) generative AI models/systems from the ground up, to use internally in production for AI-enabled processes and/or AI-enabled customer products and services.
- Build and Sell: Develop your own (proprietary) generative AI models/systems from the ground up, to make available commercially to other businesses/organizations.
- Build and Open-Source: Develop your own generative AI models/systems from the ground up, to open source to other businesses/organizations.
- Resell: Partner on or gain access to third-party generative AI models/systems to resell (modified or unmodified) to other businesses/organizations.
- Use Third Party Models/Systems: Use third-party generative AI models/systems (modified or unmodified) internally in production for AI-enabled processes and/or AI-enabled customer products and services.
- Other
- We have no plans to use/develop/sell/open-source generative AI models/systems [Exclusive]

Next 2 years. Select all that apply.

- Build and Use: Develop your own (proprietary) generative AI models/systems from the ground up, to use internally in production for AI-enabled processes and/or AI-enabled customer products and services.
- Build and Sell: Develop your own (proprietary) generative AI models/systems from the ground up, to make available commercially to other businesses/organizations.
- Build and Open-Source: Develop your own generative AI models/systems from the ground up, to open source to other businesses/organizations.
- Resell: Partner on or gain access to third-party generative AI models/systems to resell (modified or unmodified) to other businesses/organizations.
- Use Third Party Models/Systems: Use third-party generative AI models/systems (modified or unmodified) internally in production for AI-enabled processes and/or AI-enabled customer products and services.
- Other
- We have no plans to use/develop/sell/open-source generative AI models/systems [Exclusive]

q14. Now, looking at the wider AI space (excluding Generative AI), what is your organization's current and future (i.e., within two years) AI adoption strategy?

Current. Select all that apply.

- Build and Use: Develop your own (proprietary) AI models/systems from the ground up, to use internally in production for AI-enabled processes and/or AI-enabled customer products and services.
- Build and Sell: Develop your own (proprietary) AI models/systems from the ground up, to make available commercially to other businesses/organizations.
- Build and Open-Source: Develop your own AI models/systems from the ground up, to open source to other businesses/organizations.
- Resell: Partner on or gain access to third-party AI models/systems to resell (modified or unmodified) to other businesses/organizations.
- Use Third Party Models/Systems: Use third-party AI models/systems (modified or unmodified) internally in production

for AI-enabled processes and/or AI-enabled customer products and services.
- Other
- We have no plans to use/develop/sell/open-source AI models/systems [Exclusive]

Next 2 years. Select all that apply.

- Build and Use: Develop your own (proprietary) AI models/systems from the ground up, to use internally in production for AI-enabled processes and/or AI-enabled customer products and services.
- Build and Sell: Develop your own (proprietary) AI models/systems from the ground up, to make available commercially to other businesses/organizations.
- Build and Open-Source: Develop your own AI models/systems from the ground up, to open source to other businesses/organizations.
- Resell: Partner on or gain access to third-party AI models/systems to resell (modified or unmodified) to other businesses/organizations.
- Use Third Party Models/Systems: Use third-party AI models/systems (modified or unmodified) internally in production for AI-enabled processes and/or AI-enabled customer products and services.
- Other
- We have no plans to use/develop/sell/open-source AI models/systems [Exclusive]

q15. What percentage (approximately) of your AI budget are you spending/intending to spend in the next years on RAI measures?

- 0%
- 1 to 5%
- 6 to 10%
- 11 to 20%
- 21 to 30%
- 31 to 40%
- 41% or more
- Unsure
- Current
- Next 2 years

q16. Which statement best describes the extent to which the CEO and/or Board of Directors currently sponsors RAI adoption in your organization and is responsible for it?

- No active/explicit involvement
- RAI acknowledged, but not yet part of broader organization strategy
- Advocates on RAI, and is involved in some RAI efforts, e.g., initiating RAI principles
- Takes full responsibility on RAI and leads efforts, e.g., setting up RAI teams, overseeing implementation of organization-wide RAI policies and specifying RAI-related C-level KPIs
- Takes full responsibility on RAI and leads efforts, while also working with ecosystem of peers/experts/regulators to evolve RAI initiatives

q17. In your organization, RAI is viewed as... Select all that apply.

- Slowing down innovation and time to market
- Not critical for our current usage of AI
- A regulatory compliance and legal issue
- An additional cost of doing business that is necessary to incur
- A core part in shaping/driving our overall AI strategy
- A competitive advantage and value driver that can also unlock new markets
- A way to improve our/clients'/users' brand reputation and AI trustworthiness
- A differentiation in attracting/retaining key talent
- A strategic tool in better ensuring AI-related revenue growth
- A strategic tool in avoiding potential losses/brand damage due to fines, cybersecurity, etc.
- A way to demonstrate social responsibility
- A way to industrialize our AI processes and improve the performance of our AI models/systems, time to market, etc.
- Necessary to ensure the safety and security of our AI models/systems
- Other (please specify)

q18. Please select which organization-wide RAI principles, guidelines, or policies are in place in your organization. Select all that apply.

- RAI principles
- RAI guidelines
- RAI policies
- Only team-level RAI principles, guidelines or policies
- Other
- None [Exclusive]

q19. Which statement best describes your current RAI-specific governance operating model?

- None/Use of existing governance operating model with no/limited RAI focus
- RAI-specific governance framework designed, following gap analysis of existing governance model against AI risks
- RAI-specific framework translated into organization-wide operating model with RAI-specific roles and responsibilities
- RAI-specific governance operating model fully operationalized across the organization, with clear RAI roles and responsibilities, processes, reporting, tracking, KPIs and incentives for front-line personnel
- Fully-operationalized RAI-specific governance model that is extended beyond the organization to incorporate partners in supply chain/ecosystem to ensure overarching responsibility with respect to our AI models/systems

q20. Which of the following dedicated RAI roles and structures are in place in your organization (either individually, or in collaboration) to take responsibility for RAI efforts? Select all that apply.

- Cross-functional Governance Committee/Board: Consisting of multiple C-suite members, created to define and implement company RAI policies, processes, etc.
- Cross-discipline RAI team that collaborates to build/test approaches to help ensure that AI models/systems are designed and used responsibly

- AI ethics board
- Dedicated RAI role created within the C-Suite (e.g., Chief Responsible AI/AI Ethics Officer)
- RAI responsibilities are centralized with an existing, non-RAI-specific C-level role (e.g., CDO, CTO, CLO, CCO)
- Some/all RAI responsibility is decentralized across business units and functions
- Mechanism in place to jointly manage "responsibility" with external partners in the AI value chain
- Other (please specify)
- None [Exclusive]

q21. Which statement best describes the current RAI risk management framework in place at your organization?

- We don't currently have a risk management framework.
- We use our existing risk, non-AI-specific risk management framework.
- We conducted a gap analysis of the existing risk management processes completed against all major AI risks, but have yet to define an RAI risk management strategy and/or framework.
- We have defined/rolling out an RAI risk management strategy and an organization-wide RAI risk management framework.
- We have a fully operationalized RAI risk management framework across the organization, leveraging new risk management practices, AI lifecycle checkpoints and responsible parties.
- We have a fully operationalized RAI risk management framework across the organization and dedicate resources to ensure our approach stays up to date in light of advances in AI, new RAI risk management best practices and policy developments.

q22. Which statement best describes the process by which you identify risks in the development and/or use of your AI models/systems?

- We currently have no approach to identify AI risks beyond performance metrics (e.g., accuracy) and are not aware of the risks.
- We are aware of the applicable risk(s) but there is no structured approach to identify them.
- We have defined/rolling out a structured approach to risk identification, based on an analysis of the intended purpose of the AI model/system, foreseeable/known forms of misuses and potentially affected stakeholders.
- We have fully operationalized a structured approach to risk identification.
- We have fully operationalized a structured approach to risk identification, with dedicated resources in place to proactively anticipate new risks.

q23. Which statement best describes your organization's current RAI risk mitigation processes during development of your AI models/systems?

- No processes or techniques used to mitigate identified RAI risks.
- Ad-hoc mitigation techniques, typically led by data scientists/RAI practitioners at project level.
- Structured integration of some mitigation techniques into existing AI development processes.
- Integration of a comprehensive set of mitigation techniques into existing AI development processes with clear stage gates, testing/evaluation procedures and multi-disciplinary reviews.
- Integration of a comprehensive set of mitigation processes and techniques into existing AI development processes with clear stage gates, testing/evaluation procedures, multi-disciplinary reviews, and proactive research on new advances in risk mitigation.

q24. Which statement best describes how your organization applies the following measures in your current procurement evaluation processes for third-party providers when acquiring, licensing, or using AI models/systems.
Scale:

1. Not applied
2. Assess – We are assessing this measure but haven't used it yet
3. Ad-Hoc – We have used this measure on an ad-hoc basis with our models/systems
4. Rolling out – We have defined requirements and are rolling this out for all our relevant models/systems
5. Fully operationalized – We require this measure to be implemented for all our relevant AI models/systems

Subquestions:

a. Key metrics for RAI dimensions like performance, reliability, and fairness
b. Provider's RAI risk identification and evaluation processes
c. Provider's documentation of model limitations, instructions for use, interfaces and tailored training
d. Demonstration of alignment of model/system with emerging regulation of AI in relevant jurisdictions
e. Service agreements that encompass redress channels and remediation of RAI-related risks
f. Contractual agreements to ensure all RAI-related responsibilities have been agreed upon and are fully documented
g. Checks to ensure that the data used to train models was legally obtained

q25. Which statement best describes the RAI monitoring and control processes that your organization has in place after you've deployed an AI model/system or after you've given access to your model/system to an external party?

- No specific monitoring and control process in place
- Some general (non-RAI specific) monitoring by MLOps and/or similar teams within the organization
- Centralized plan for RAI monitoring and control defined/rolling out to monitor and identify risks in production, with defined cross-domain responsibilities, roles, and processes
- RAI monitoring and control system fully operationalized to monitor and identify risks in production, with defined cross-domain responsibilities, roles, and processes

- Monitoring and control systems fully operationalized and optimized to proactively and continuously identify new RAI risks

q26. What support do you provide to purchasers/users of your AI models, systems, and services (if any) to help them meet their RAI priorities or obligations (e.g., with respect to fairness, transparency, or compliance)? Select all that apply.

- Mechanisms to provide support to purchasers/users with select RAI priorities and obligations, on request
- Comprehensive guidance and support for purchasers to meet their RAI priorities and obligations, as standard
- We have built our product to anticipate and meet RAI priorities and obligations and provide additional support where necessary
- Other
- No specific RAI support [Exclusive]

q27. Which statement best describes how your organization adopts cybersecurity in your AI risk management practices?

- We are not aware of any AI-related cybersecurity risks and do not have in place any related measures.
- We are aware of the AI-related cybersecurity risks but do not have in place any related cybersecurity measures.
- We use our existing cybersecurity policies in an ad-hoc manner to address AI-related incidents.
- We have a clearly outlined AI-specific cybersecurity incident response plan for containing, investigating, and rectifying AI-related events.
- We have clear AI-specific cybersecurity incident response plan with vulnerability management measures for model audits, adversarial testing, etc., within our AI risk-management framework.

q28. Which statement best describes how your organization applies the following measures to mitigate human interaction risks.

Scale:

1. Not applied
2. Assess – We are assessing this measure but haven't used it yet
3. Ad-Hoc – We have used this measure on an ad-hoc basis with our models/systems
4. Rolling out – We have defined requirements and are rolling this out for all our relevant models/systems
5. Fully operationalized – We require this measure to be implemented for all our relevant AI models/systems

Subquestions:

a. Implementation of technical safeguards on top of the base model (e.g., filters) to decrease non-compliant behavior or misuse potential
b. Use of human-centered design techniques to specifically improve user understanding e.g., interface design, etc.
c. Monitor usage patterns for anomalous activity that may indicate misuse or attempts to game the system
d. Case/Role-specific training of users, and provision of information about the limitations of the AI model/system
e. Use of codes of conduct or terms of service to restrict misuse
f. Testing for potential misuse of AI model/system during evaluation

q29. Which statement best describes how your organization applies the following measures to ensure sufficient data quality, both during training and with respect to the input data used during deployment.

Scale:

1. Not applied
2. Assess – We are assessing this measure but haven't used it yet
3. Ad-Hoc – We have used this measure on an ad-hoc basis with our models/systems
4. Rolling out – We have defined requirements and are rolling this out for all our relevant models/systems
5. Fully operationalized – We require this measure to be implemented for all our relevant AI models/systems

Subquestions:

a. Data collection and preparation include assessment of the completeness, uniqueness, consistency, and accuracy of the data
b. Remediation plans for and documentation of datasets with shortcomings
c. Checks to ensure that the data is representative with respect to the demographic setting within which the final model/system is used
d. Regular data audits and updates to ensure the relevancy of the data
e. Checks to ensure that the data complies with all relevant laws and regulations and is used with consent where applicable
f. Process for dataset documentation and traceability throughout the AI lifecycle

q30. Which statement best describes how your organization applies the following measures to increase AI model/system fairness.

Scale:

1. Not applied
2. Assess – We are assessing this measure but haven't used it yet
3. Ad-Hoc – We have used this measure on an ad-hoc basis with our models/systems
4. Rolling out – We have defined requirements and are rolling this out for all our relevant models/systems
5. Fully operationalized – We require this measure to be implemented for all our relevant AI models/systems

Subquestions:

a. Collection of representative data based on the anticipated user demographics
b. Making methodology and data sources accessible to third-parties (auditors/general public) for independent oversight
c. Involvement of diverse stakeholders in model development and/or review process
d. Assessment of performance across different demographic groups
e. Use of technical bias mitigation techniques during model development

q31. Which statement best describes how your organization applies the following measures to measure and minimize AI model/system environmental footprint.

Scale:

1. Not applied
2. Assess – We are assessing this measure but haven't used it yet
3. Ad-Hoc – We have used this measure on an ad-hoc basis with our models/systems
4. Rolling out – We have defined requirements and are rolling this out for all our relevant models/systems
5. Fully operationalized – We require this measure to be implemented for all our relevant AI models/systems

Subquestions:

a. Measurement of environmental footprint of AI models/systems
b. Provision of carbon impact statement for AI models/systems
c. Technical measures to minimize environmental impact during AI development (e.g., by implementing power-efficient coding practices or using eco-friendly hardware)
d. Technical measures to minimize environmental impact during use/deployment (e.g., by using energy-efficient infrastructure or setting energy-efficient default hyperparameters)
e. Carbon reduction strategies at the organization level (e.g., carbon offsetting or use of renewable energy)

q32. Which statement best describes how your organization applies the following measures to validate and ensure AI model/system reliability.

Scale:

1. Not applied
2. Assess – We are assessing this measure but haven't used it yet
3. Ad-Hoc – We have used this measure on an ad-hoc basis with our models/systems
4. Rolling out – We have defined requirements and are rolling this out for all our relevant models/systems
5. Fully operationalized – We require this measure to be implemented for all our relevant AI models/systems

Subquestions:

a. Mitigation measures for model errors and to handle low confidence outputs
b. Failover plans or other measures to ensure the system's/model's availability
c. Evaluation of models/systems for vulnerabilities or harmful behavior (i.e., red teaming)
d. Measures to prevent adversarial attacks
e. Confidence scoring for model outputs
f. Comprehensive test cases that cover a wide range of scenarios and metrics

q33. Which statement best describes how your organization applies the following measures to increase model/system transparency.

Scale:

1. Not applied
2. Assess – We are assessing this measure but haven't used it yet
3. Ad-Hoc – We have used this measure on an ad-hoc basis with our models/systems
4. Rolling out – We have defined requirements and are rolling this out for all our relevant models/systems
5. Fully operationalized – We require this measure to be implemented for all our relevant AI models/systems

Subquestions:

a. Model explainability/interpretability tools (e.g., saliency maps) to elucidate model decisions
b. Documentation of the development process, detailing algorithm design choices, data sources, intended use cases, and limitations
c. Prioritization of simpler models where high interpretability is crucial, even if it sacrifices some performance
d. Training programs for stakeholders (incl. users) covering the intended use cases and limitations of model

q34. Which statement best describes how your organization applies the following measures to increase model/system accountability.

Scale:

1. Not applied
2. Assess – We are assessing this measure but haven't used it yet
3. Ad-Hoc – We have used this measure on an ad-hoc basis with our models/systems
4. Rolling out – We have defined requirements and are rolling this out for all our relevant models/systems
5. Fully operationalized – We require this measure to be implemented for all our relevant AI models/systems

Subquestions:

a. Logging and traceability mechanisms of a model's/system's outputs
b. Redress mechanisms for stakeholders
c. Negative incidence tracking and reporting
d. Publication of model/system documentation detailing algorithm design choices, data sources, intended use cases, and limitations
e. Third-party audits of models/systems for RAI considerations
f. Version control of models/systems to track changes, updates, or modifications
g. Regular model/system reviews to ensure they're aligned with original objectives

q35. Which statement best describes how your organization applies the following measures to improve model/system cybersecurity.

Scale:

1. Not applied
2. Assess – We are assessing this measure but haven't used it yet
3. Ad-Hoc – We have used this measure on an ad-hoc basis with our models/systems
4. Rolling out – We have defined requirements and are rolling this out for all our relevant models/systems
5. Fully operationalized – We require this measure to be implemented for all our relevant AI models/systems

Subquestions:

a. Basic cybersecurity hygiene practices (e.g., multi-factor authentication, access controls, and employee training)
b. Dedicated AI cybersecurity team and/or personnel trained specifically for AI-specific cybersecurity
c. Technical AI-specific cybersecurity checks and measures, e.g., adversarial testing, vulnerability assessments, and data security measures
d. Resources dedicated to research and monitoring of evolving AI-specific cybersecurity risks and integration in existing cybersecurity processes

e. Vetting and validation of cybersecurity measures of third-parties in the supply chain

q36. Which statement best reflects the current level of RAI training offered across your organization?

- No official RAI training currently available
- Ad hoc training in siloes/by project
- Guidance and high-level training provided for those holding RAI responsibilities, including end users, business owners, data scientists, etc.
- RAI training is part of all AI-related employee training. One-time role-specific RAI training provided to all employees with RAI responsibilities, including end users, business owners, data scientists, etc.
- RAI training is part of all AI-related employee training. Ongoing role-specific RAI training provided to all employees with RAI responsibilities including end users, business owners, data scientists, etc.

q37. From the following list, please select the areas where you feel your organization has sufficient talent/skills to conduct RAI activities effectively in the next 2 years.
Scale:

1. Sufficient
2. Insufficient

Subquestions:

a. RAI strategic/organizational governance
b. Compliance and legal RAI landscape
c. Technical AI risk mitigation
d. Human-centered design
e. AI-specific cybersecurity

q38. Please indicate how much you agree with the following statements applicability to your organization in the next 5 years.
Scale:

1. Strongly disagree
2. Disagree
3. Neither agree nor disagree
4. Agree
5. Strongly agree
6. Not sure

Subquestions:

a. My organization will be subject to liability laws and directives due to our AI adoption.
b. My organization is or will be subject to cybersecurity laws due to our AI adoption.
c. My organization is or will be subject to data protection and privacy laws due to our AI adoption.
d. My organization is or will be subject to consumer safety & protection laws due to our AI adoption.
e. My organization expects to be subject to future generative AI law (e.g., regarding truthfulness/misinformation) due to our AI adoption.
f. My organization contributes to RAI best practice and standardization efforts.

q39. What are the greatest barriers to the use/development of generative AI by your organization? Rank the top three, where 1 = greatest barrier.

(1) Lack of a clear vision, roadmap, infrastructure, etc.
(2) Lack of C-Suite/board support
(3) Lack of clear ROI/use cases associated with of generative AI
(4) Inability to move beyond early pilots and proof of concepts to scaling across the organization
(5) Shortage of, or difficult access to, talent with generative AI skills
(6) Data and privacy concerns
(7) Risks associated with the use of generative AI (e.g., hallucination/factually incorrect output/transparency and explainability, security and bias or misuse related concerns etc.)
(8) Lack of budget
(9) Challenges due to current/future regulation of generative AI
(10) Lack of trustworthy generative AI providers
(11) Other (please specify)

q40. Which of the following measures do you have in place to mitigate risks when using generative AI models from a third-party provider? Select all that apply.

- Provider Selection: Selection of model/system provider that prioritizes and supports RAI considerations
- Evaluation: RAI-risk-focused testing of models (e.g., human feedback, red teaming for toxic/offensive output, adversarial testing)
- Infrastructure: Running generative AI models/systems on infrastructure with appropriate security and privacy protocols, e.g., to prevent data leakages or system breaches
- Application: Implementation of additional technical safeguards on top of third-party model/system to decrease non-compliant behavior or misuse potential (e.g., addition of safety packages or filters to prevent toxic/offensive content, hallucinations, malicious prompts, prompt injections, or PII & sensitive data leaks)
- End-User: Any measure targeted at end-users, e.g., use of codes of conduct or terms of service and training to restrict misuse
- Monitoring, Control & Observability: Monitoring/analysis of prompts, responses and model behavior to mitigate issues and improve performance
- Other
- None [Exclusive]

q41. Which of the following measures do you have in place to mitigate risks when developing generative AI models/systems? Select all that apply.

- Infrastructure: Running generative AI systems on infrastructure with appropriate security and privacy protocols, e.g., to prevent data leakage or system breaches
- Model: Technical risk mitigation approaches at the model level, e.g., reinforcement learning from human feedback or fine-tuning
- Evaluation: RAI-risk-focused testing of models (e.g., evaluator GAIs, red teaming for toxic/offensive output)

- Application: Implementation of additional technical safeguards on top of base model to decrease non-compliant behavior or misuse potential (e.g., addition of safety packages and filters to prevent toxic/offensive content, hallucinations, malicious prompts, prompt injections, or PII & sensitive data leaks)
- End-User: Any measure targeted at end-users, e.g., use of codes of conduct or terms of service and training to restrict misuse
- Post-Deployment monitoring, control & observability: Monitoring/analysis of prompts, responses and model behavior to mitigate issues and improve performance
- Other
- None [Exclusive]

q42. Please indicate how much you agree with the following statements.

Scale:

1. Strongly disagree
2. Disagree
3. Neither agree nor disagree
4. Agree
5. Strongly agree
6. Not sure

Subquestions:

a. Companies that develop foundation models will be responsible for the mitigation of all associated risks, rather than organizations using these models/systems.
b. The new and evolving business risks associated with generative AI mean that RAI will be essential for my organization to unlock and maintain its value over time.
c. Employees/end-users will play a pivotal role in the identification and mitigation of risks (e.g., hallucinations or cybersecurity and IP/data breaches).
d. I believe that generative AI presents enough of a threat that globally agreed generative AI governance is required.
e. Using open-source generative AI models is safer than using proprietary generative AI models because there's added public scrutiny and we can make direct modifications to the model, if necessary.

C Additional Results

C.0.1 Regression Results. Tab. 6 presents the OLS regression analysis results for organizational maturity. The results indicate that the diversity and non-discrimination risks (Q9-1) and privacy & data governance risks (Q9-2) have negative associations with organizational maturity, with coefficients of −2.78 and −2.27, respectively, and are statistically significant. Also, environmental risk (Q9-8) and brand & reputational risk have positive associations with organizational maturity, with coefficients of 2.81 and 2.17, respectively and are statistically significant. Other risk factors, such as security risks (Q9-4) and human interaction risks (Q9-5), show no statistically significant relationship with organizational maturity.

Tab. 7 shows the relationship between the total number of risks identified by organizations (sum-Q1-Q13) and their organizational maturity. The coefficient for sum-Q1-Q13 is 0.1892, with a p-value of 0.400, indicating that the total number of identified risks does not have a statistically significant impact on organizational maturity. This suggests that simply identifying more risks does not necessarily correlate with higher or lower organizational maturity levels.

Tab. 8 presents the regression results for the impact of the organization's expectation to be subject to AI-specific regulations or legal liabilities (Q10) on organizational maturity. The results reveal a negative association with a coefficient of −2.1318 and a p-value of 0.000, indicating statistical significance.

Tab. 9 summarizes the regression analysis of various AI-related risk factors on operational maturity. The results show that diversity and non-discrimination risks (Q9-1) have a positive and statistically significant impact on operational maturity with a coefficient of 2.8510. Conversely, privacy & data governance risks (Q9-2) and accountability risks (Q9-12) have negative associations, with coefficients of −0.2123 and −3.9899, respectively. This suggests that while addressing diversity issues may enhance operational maturity, privacy and accountability concerns could pose challenges to operational maturity.

Tab. 10 details the impact of the total number of identified risks (sum-Q1-Q13) on operational maturity. The coefficient is −0.6454 with a p-value of 0.034, indicating a significant negative relationship. This suggests that organizations identifying a higher number of risks may experience lower operational maturity, possibly due to the complexity and resource demands of managing multiple risks simultaneously.

The analysis presented in Tab. 11 examines the effect of organizations' expectations of future AI regulation (Q10) on operational maturity. The results show a negative relationship with a coefficient of −2.3030 and a p-value of 0.000, suggesting a statistically significant impact on operational maturity. This may reflect concerns over the readiness and adaptability of operational processes in the face of impending regulatory demands.

Variable	Coefficient	Standard Error	t-Statistic	p-Value	95% Confidence Interval
Constant	69.2131	1.267	54.614	0.000	[66.726, 71.700]
q9_1	-2.7819	0.955	-2.912	0.004	[-4.657, -0.907]
q9_2	-2.2698	0.919	-2.469	0.014	[-4.074, -0.465]
q9_3	0.6507	0.927	0.702	0.483	[-1.169, 2.471]
q9_4	0.1571	0.911	0.172	0.863	[-1.631, 1.945]
q9_5	0.3374	0.966	0.349	0.727	[-1.558, 2.232]
q9_6	1.3797	0.983	1.403	0.161	[-0.549, 3.309]
q9_7	0.3361	0.937	0.359	0.720	[-1.502, 2.174]
q9_8	2.8151	1.005	2.800	0.005	[0.842, 4.788]
q9_9	0.4066	0.960	0.424	0.672	[-1.477, 2.290]
q9_10	2.1711	1.006	2.158	0.031	[0.196, 4.146]
q9_11	-0.5122	1.027	-0.499	0.618	[-2.528, 1.504]
q9_12	0.2019	1.030	0.196	0.845	[-1.819, 2.222]
q9_13	0.3036	1.383	0.220	0.826	[-2.410, 3.017]

Table 6: OLS regression results for organizational maturity.

Variable	Coefficient	Standard Error	t-Statistic	p-Value	95% Confidence Interval
Constant	69.1004	1.090	63.420	0.000	[66.962, 71.239]
sum_q1_q13	0.1892	0.225	0.842	0.400	[-0.252, 0.630]

Table 7: OLS regression analysis results for organizational maturity.

Variable	Coefficient	Standard Error	t-Statistic	p-Value	95% Confidence Interval
Constant	74.9482	0.998	75.102	0.000	[72.990, 76.907]
q10	-2.1318	0.386	-5.527	0.000	[-2.889, -1.375]

Table 8: OLS regression analysis results for organizational maturity.

Variable	Coefficient	Standard Error	t-Statistic	p-Value	95% Confidence Interval
Constant	38.3211	1.726	22.202	0.000	[34.934, 41.708]
q9_1	2.8510	1.301	2.191	0.029	[0.298, 5.404]
q9_2	-0.2123	1.252	-0.170	0.865	[-2.670, 2.245]
q9_3	0.1327	1.263	0.105	0.916	[-2.346, 2.612]
q9_4	2.1627	1.241	1.743	0.082	[-0.273, 4.598]
q9_5	-3.4642	1.315	-2.634	0.009	[-6.045, -0.883]
q9_6	-0.8716	1.339	-0.651	0.515	[-3.499, 1.756]
q9_7	-0.9703	1.276	-0.761	0.447	[-3.473, 1.533]
q9_8	1.4518	1.369	1.060	0.289	[-1.235, 4.139]
q9_9	-2.9355	1.307	-2.245	0.025	[-5.501, -0.370]
q9_10	0.4525	1.370	0.330	0.741	[-2.237, 3.142]
q9_11	-0.1418	1.399	-0.101	0.919	[-2.888, 2.604]
q9_12	-3.9899	1.402	-2.845	0.005	[-6.742, -1.238]
q9_13	-0.3460	1.883	-0.184	0.854	[-4.042, 3.350]

Table 9: OLS regression analysis results for operational maturity.

Variable	Coefficient	Standard Error	t-Statistic	p-Value	95% Confidence Interval
Constant	39.4185	1.478	26.678	0.000	[36.519, 42.318]
sum_q1_q13	-0.6454	0.305	-2.119	0.034	[-1.243, -0.048]

Table 10: OLS regression analysis results for operational maturity.

Variable	Coefficient	Standard Error	t-Statistic	p-Value	95% Confidence Interval
Constant	41.9424	1.364	30.756	0.000	[39.266, 44.618]
q10	-2.3030	0.527	-4.369	0.000	[-3.337, -1.269]

Table 11: OLS regression analysis results for operational maturity.

C.0.2 Region and Industry-Specific Results. This section showcases organizational and operational RAI maturity scores for specific regions and industries.

C.0.3 Other Results.

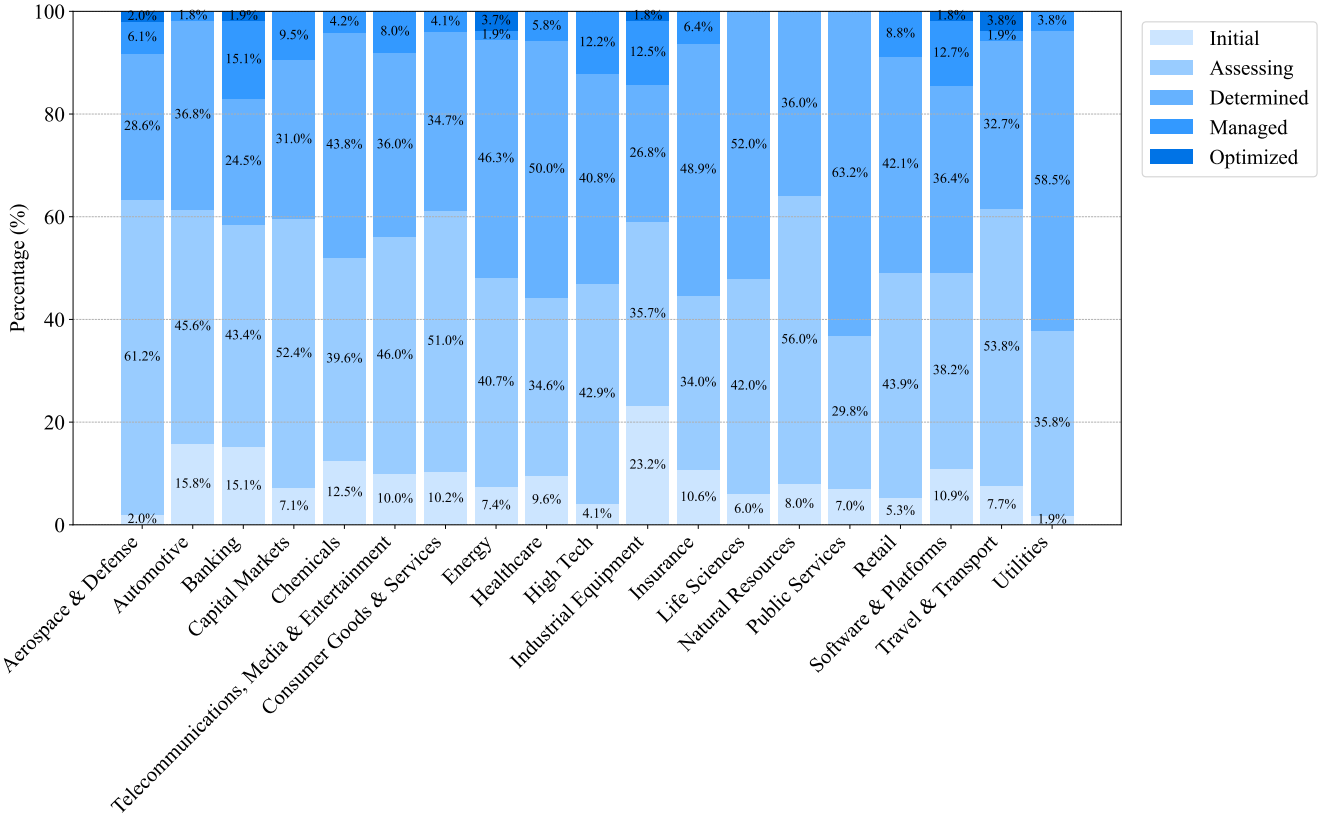

Figure 7: Operational maturity level across different industries.

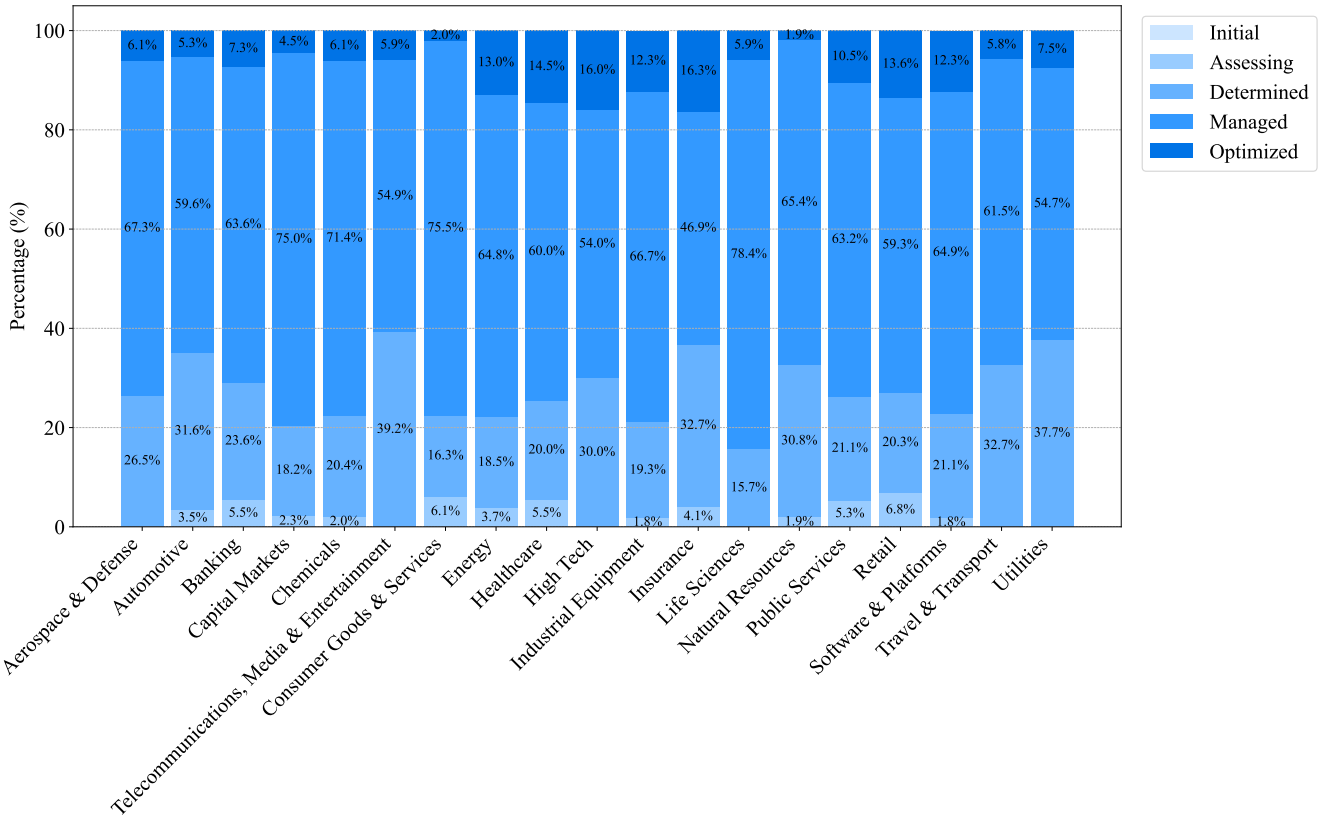

Figure 8: Organizational maturity level across different industries.

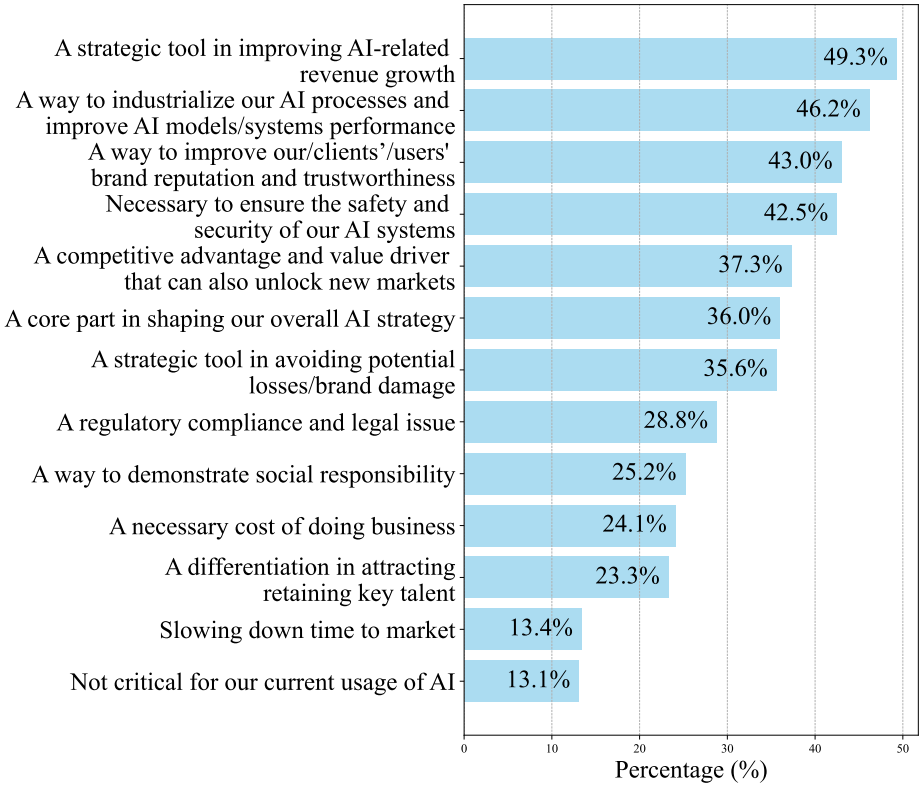

Figure 9: The way RAI is viewed within organizations.

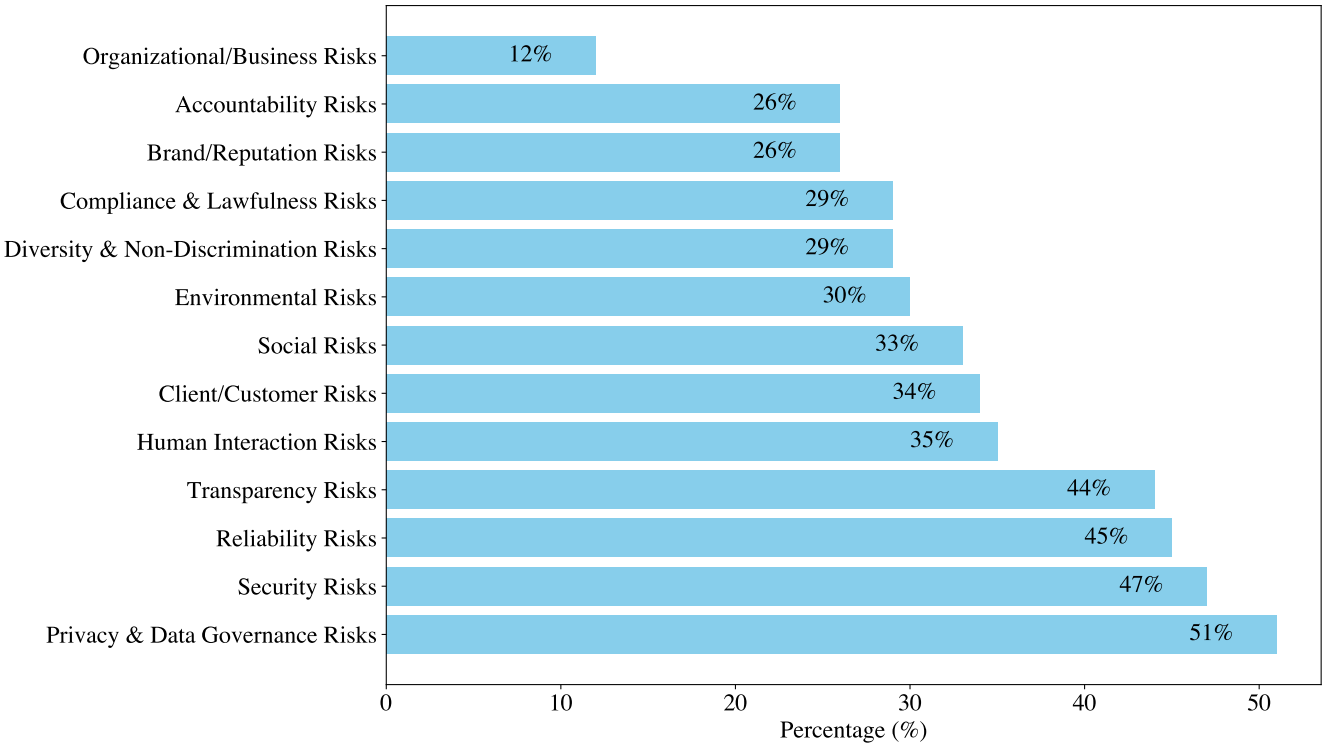

Figure 10: Risks selected by the respondents (multiple selections were possible).

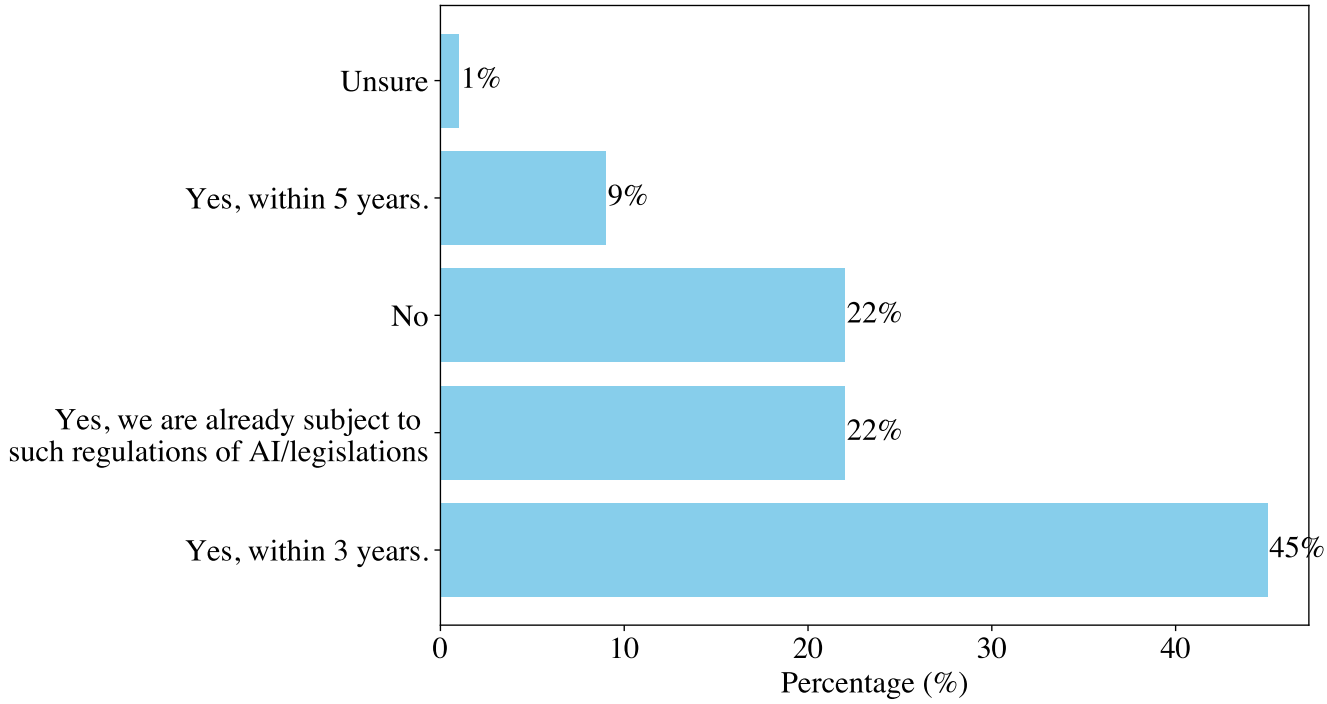

Figure 11: (Expected) regulatory exposure to specific AI regulations and legal liabilities.

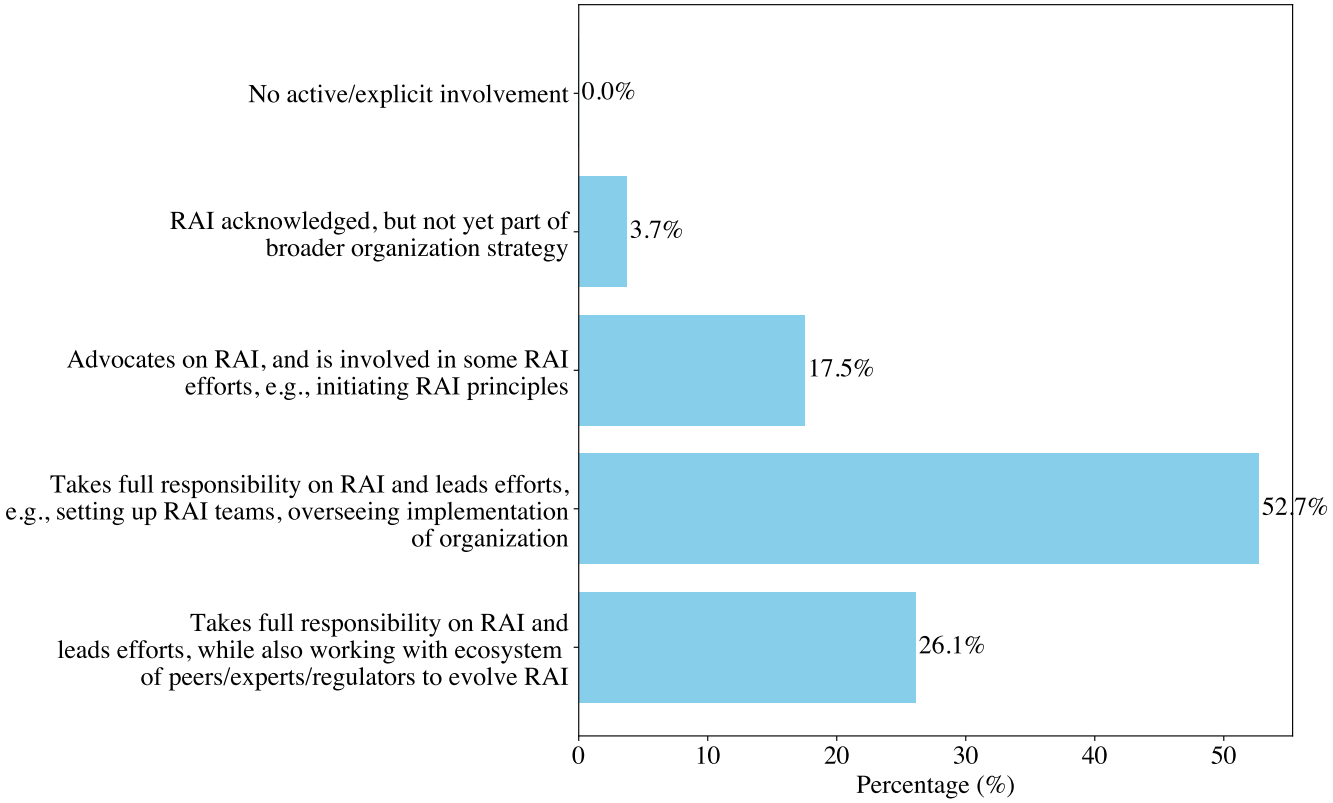

Figure 12: Involvement of CEO/Board of Directors in RAI initiatives.

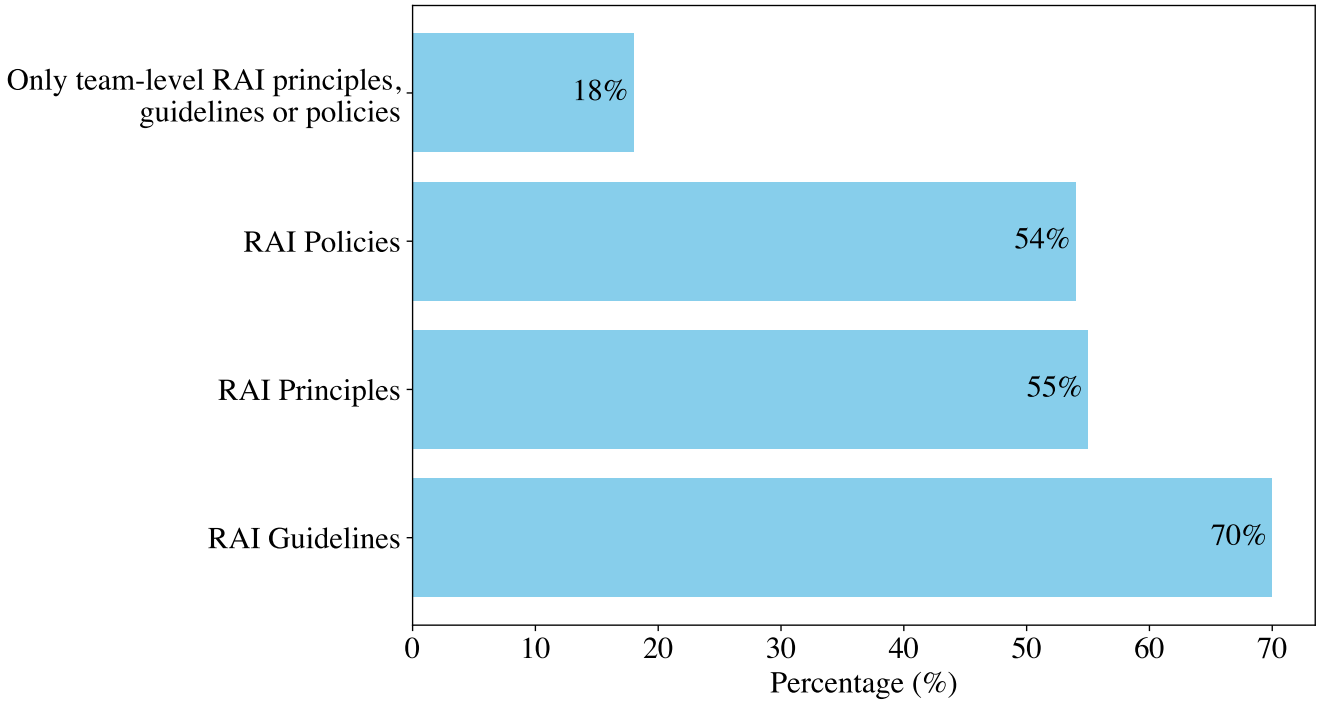

Figure 13: Indication if organization-wide RAI principles, guidelines, or policies are in place in the organization.

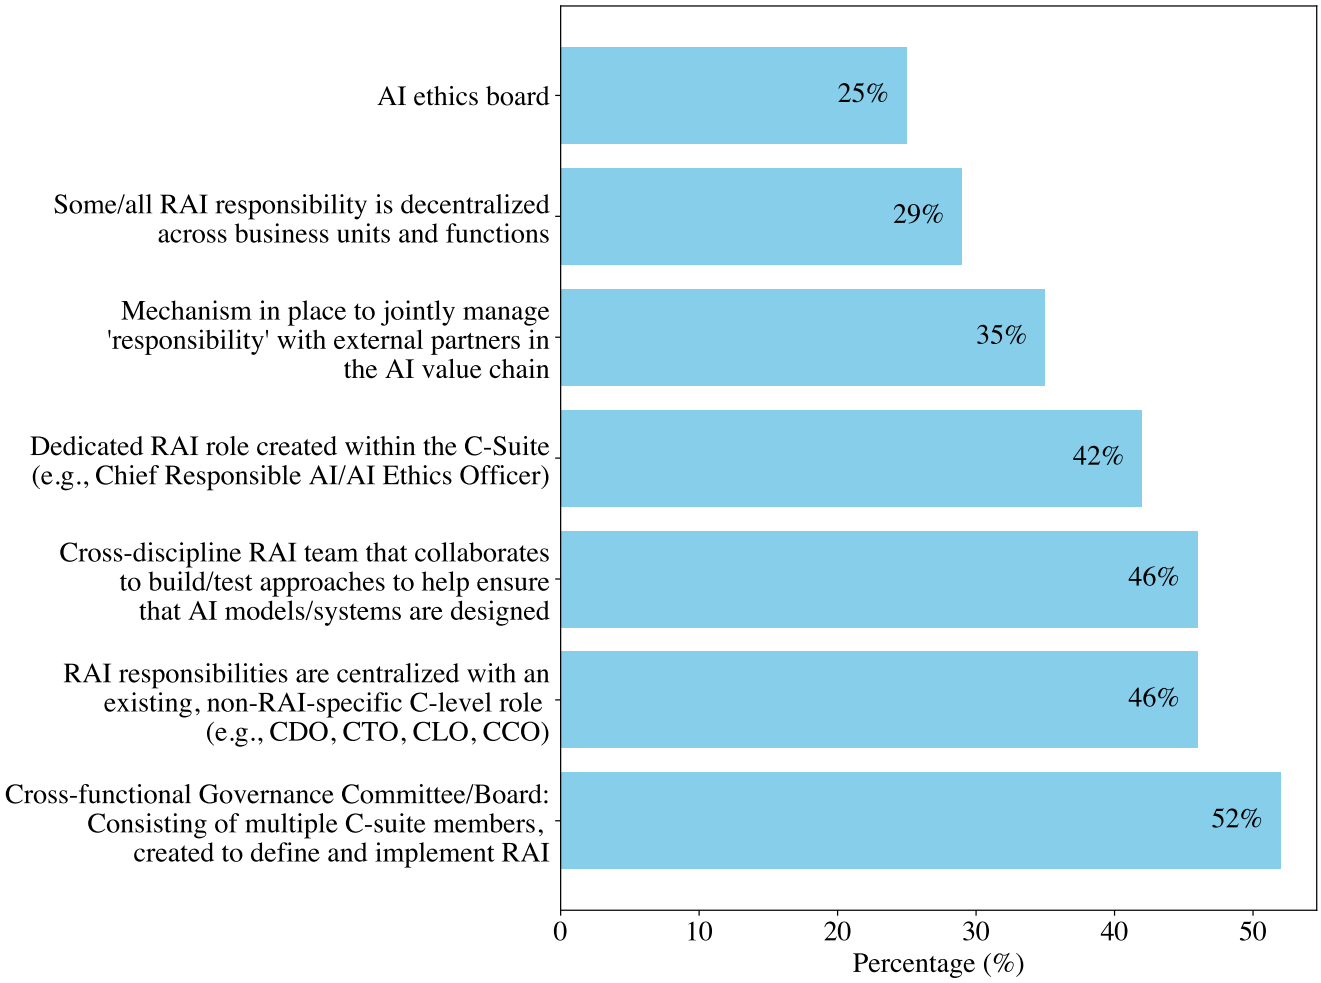

Figure 14: Dedicated RAI roles and structures that are in place in organizations.

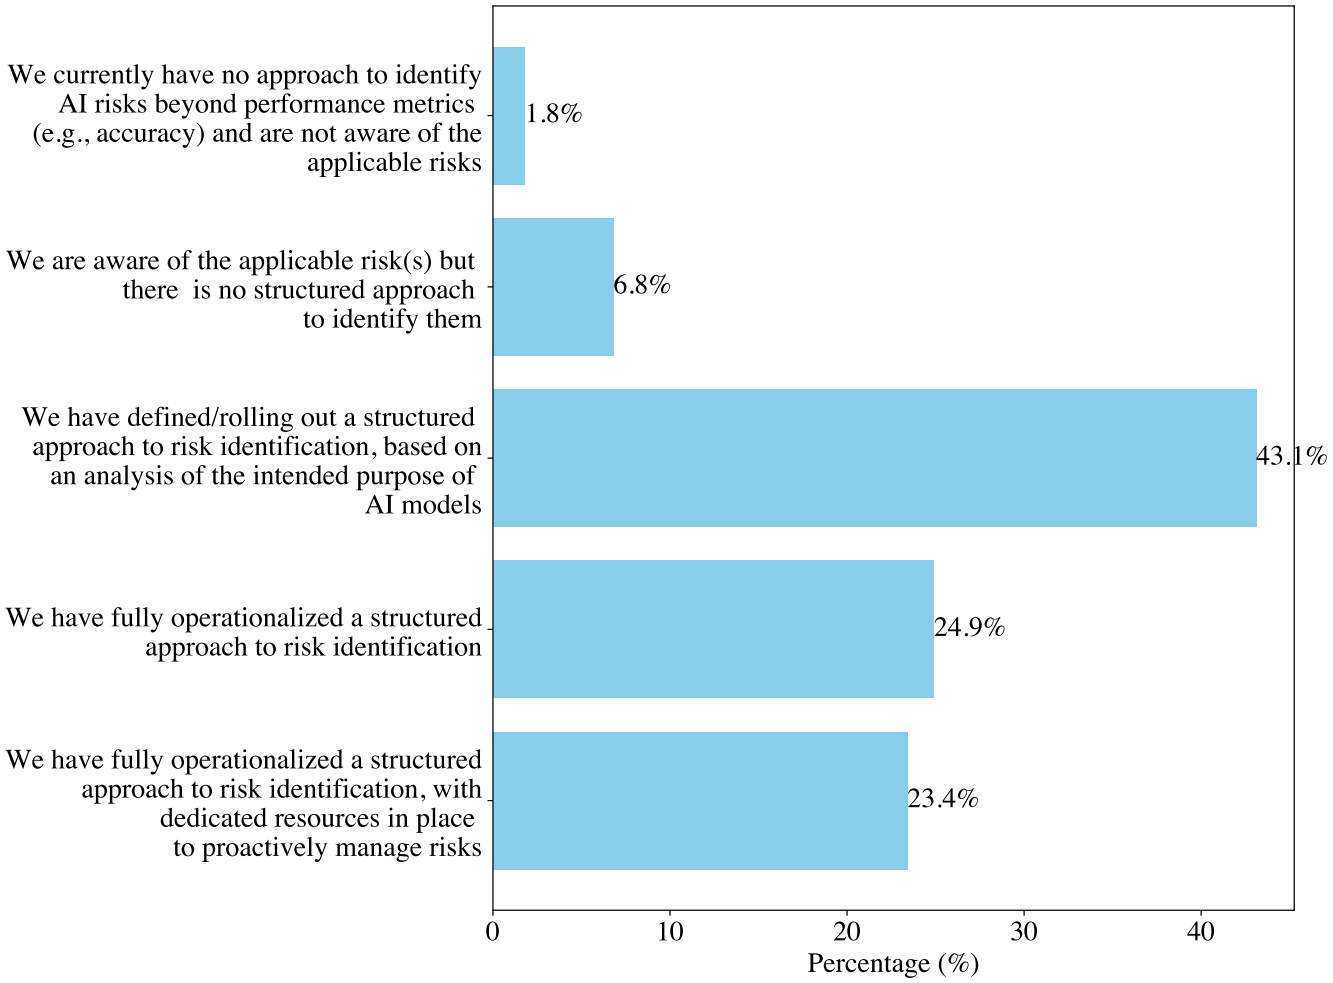

Figure 15: Risk identification processes in surveyed organizations.

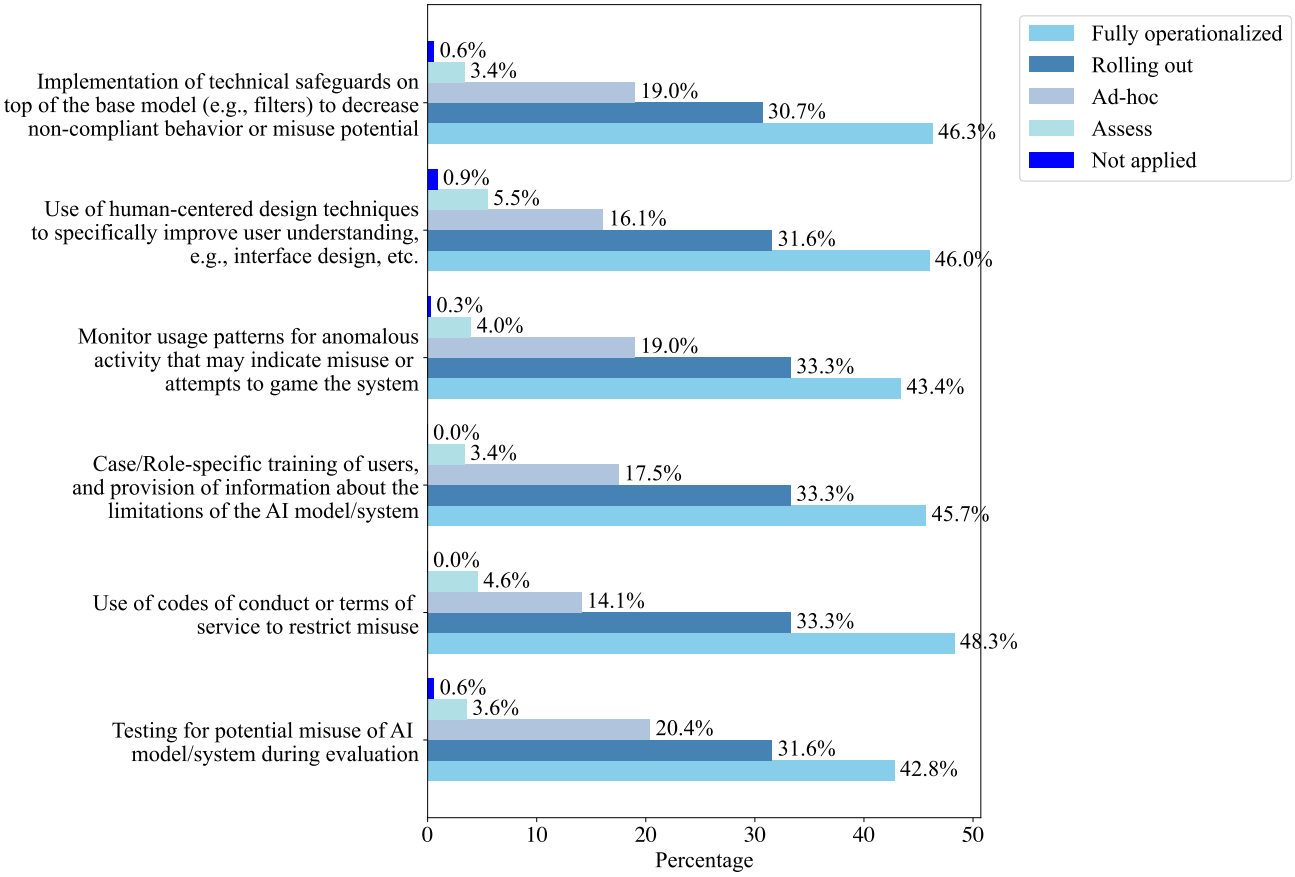

Figure 16: Implemented RAI measures to address human interaction risks.

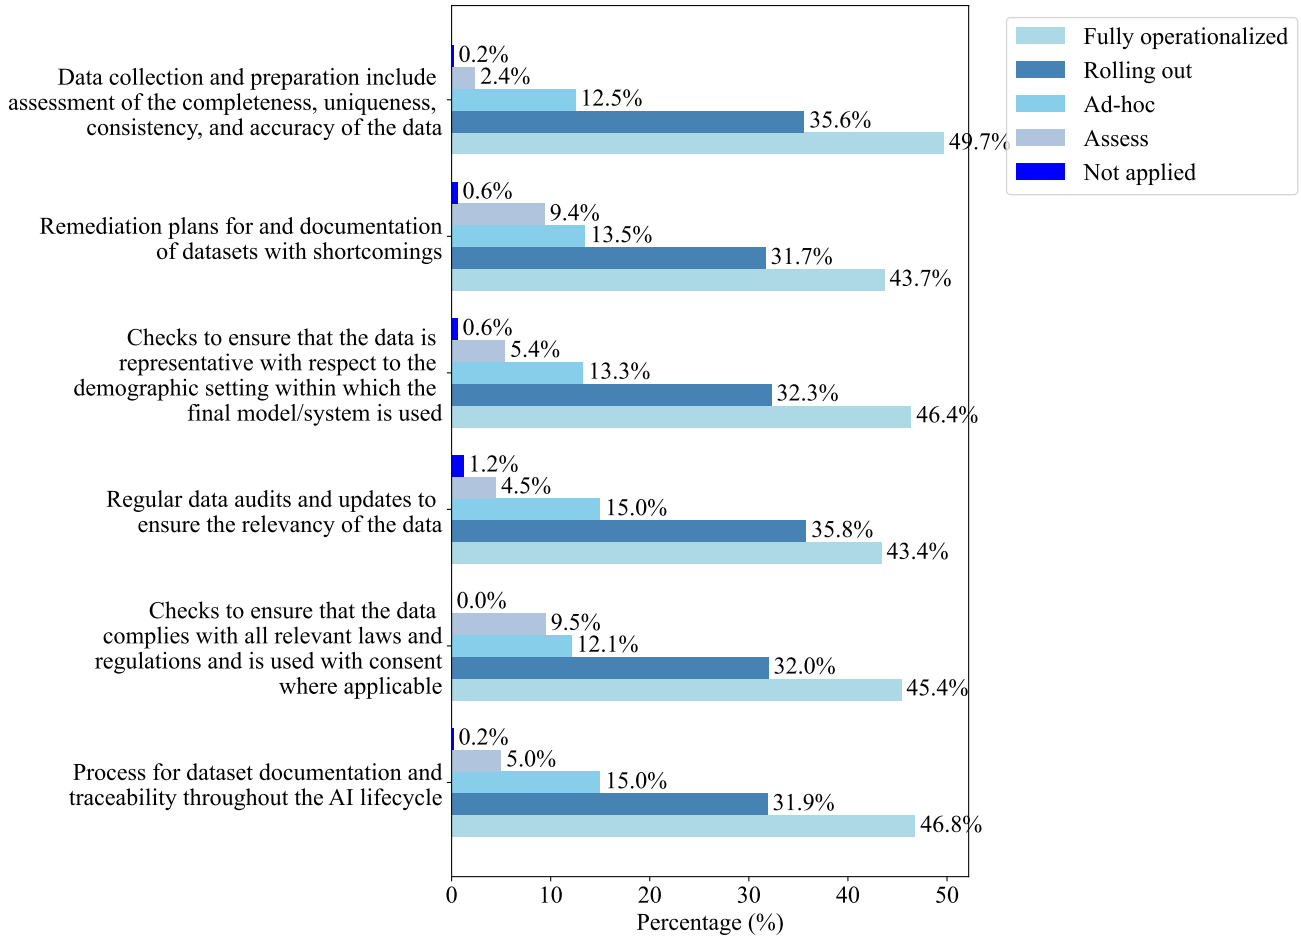

Figure 17: Implemented RAI measures to address data-related concerns.

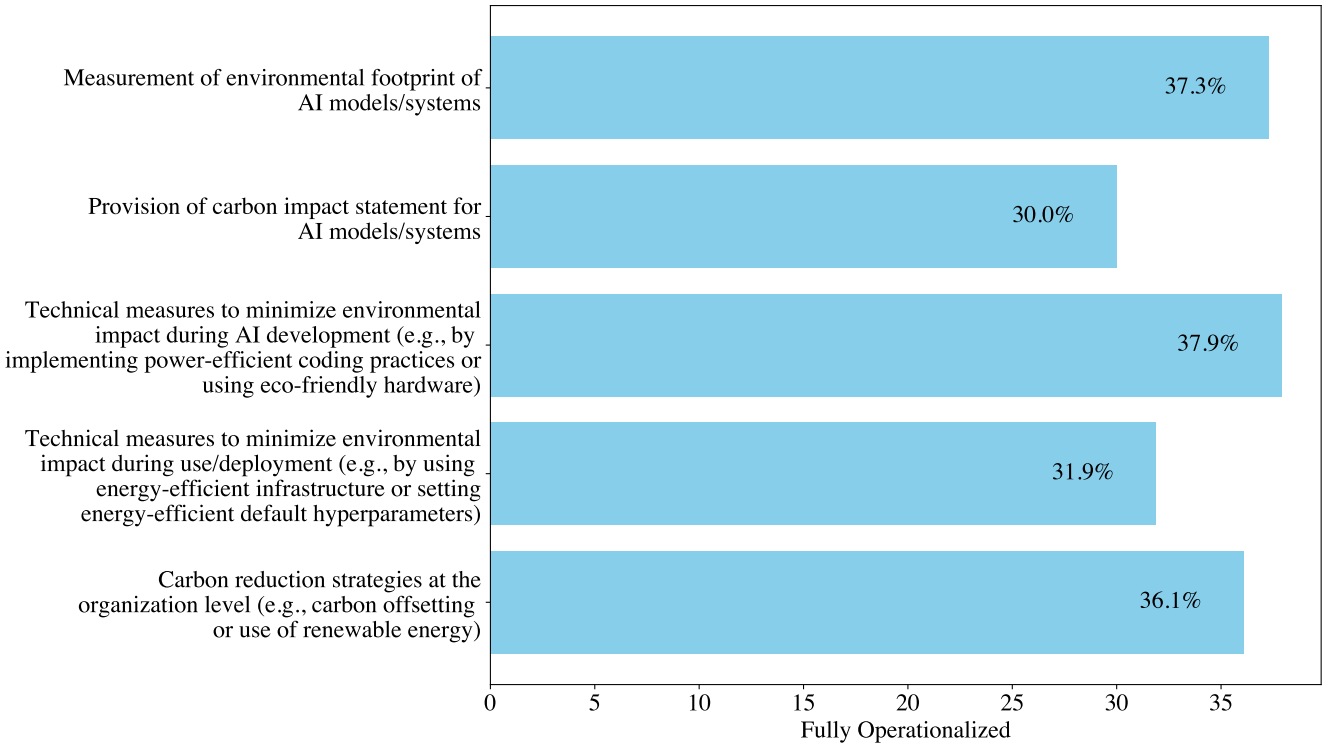

Figure 18: Implemented RAI measures to address sustainability risks.

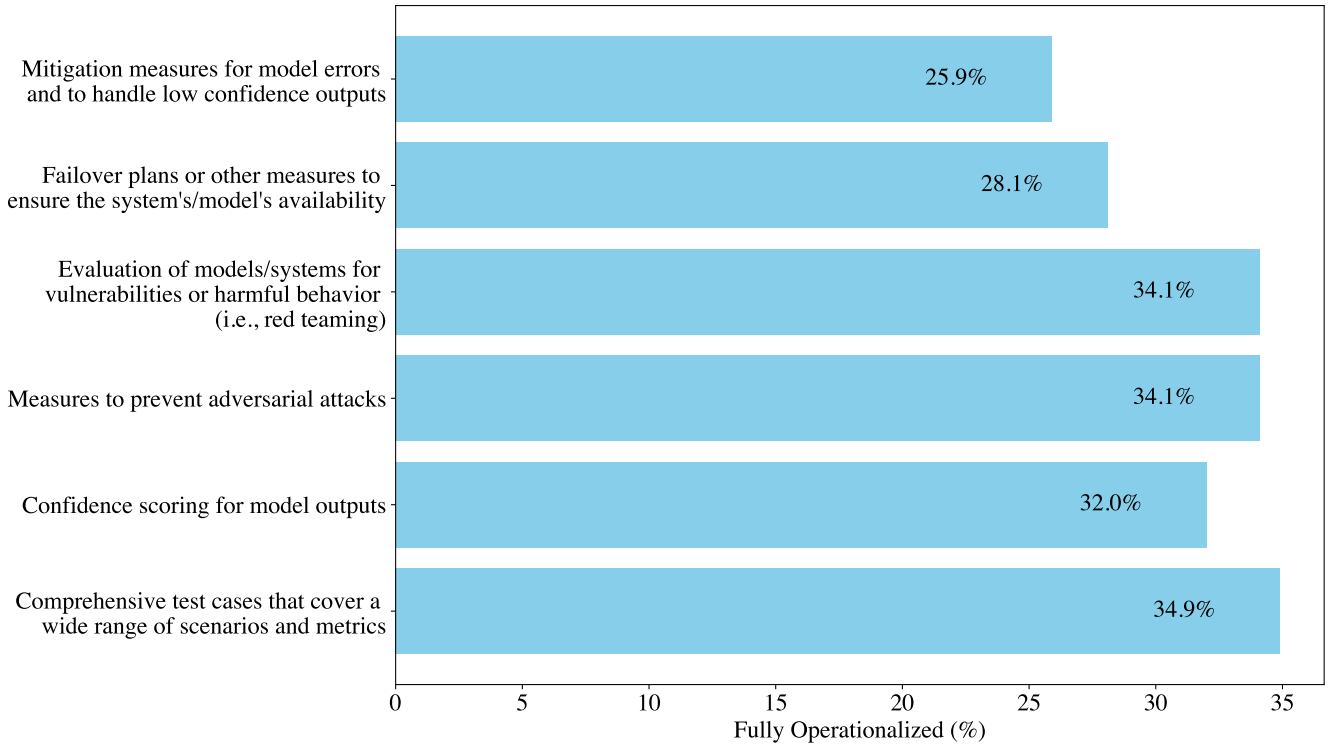

Figure 19: Implemented RAI measures to address reliability risks.

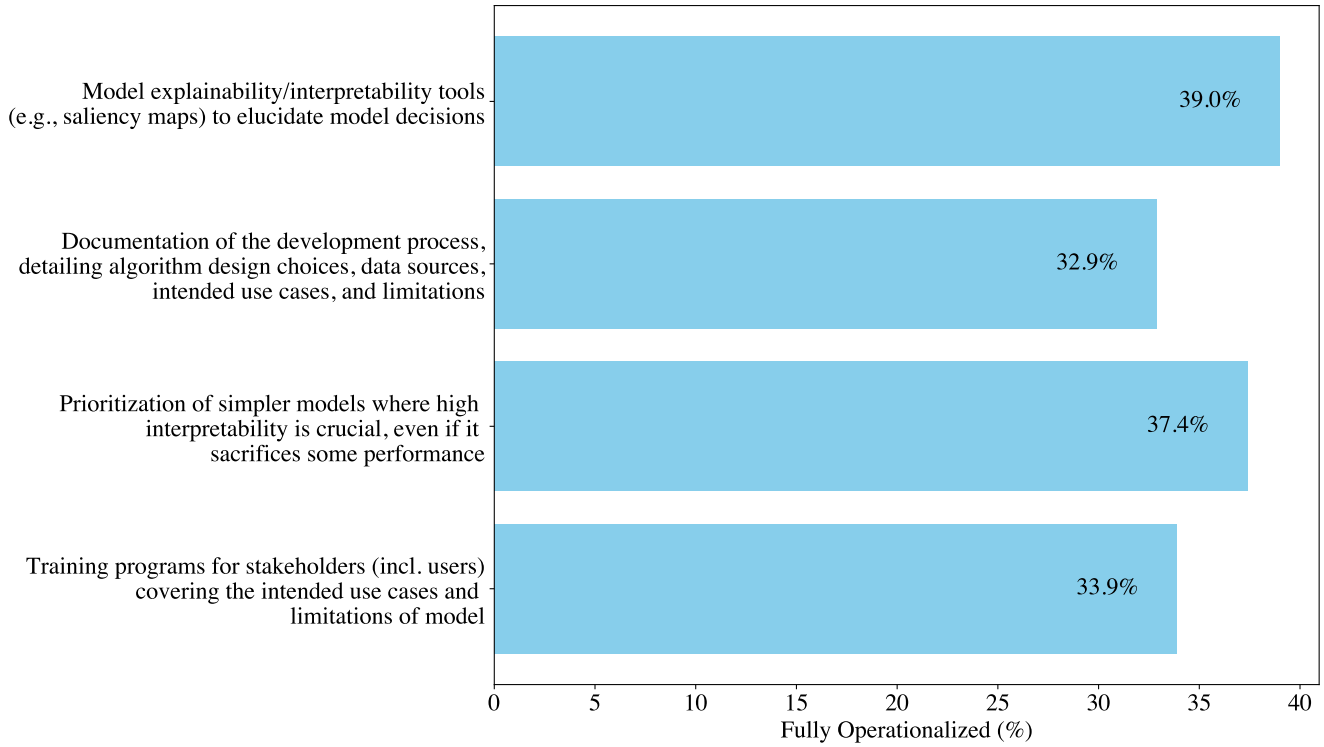

Figure 20: Implemented RAI measures to address transparency concerns.

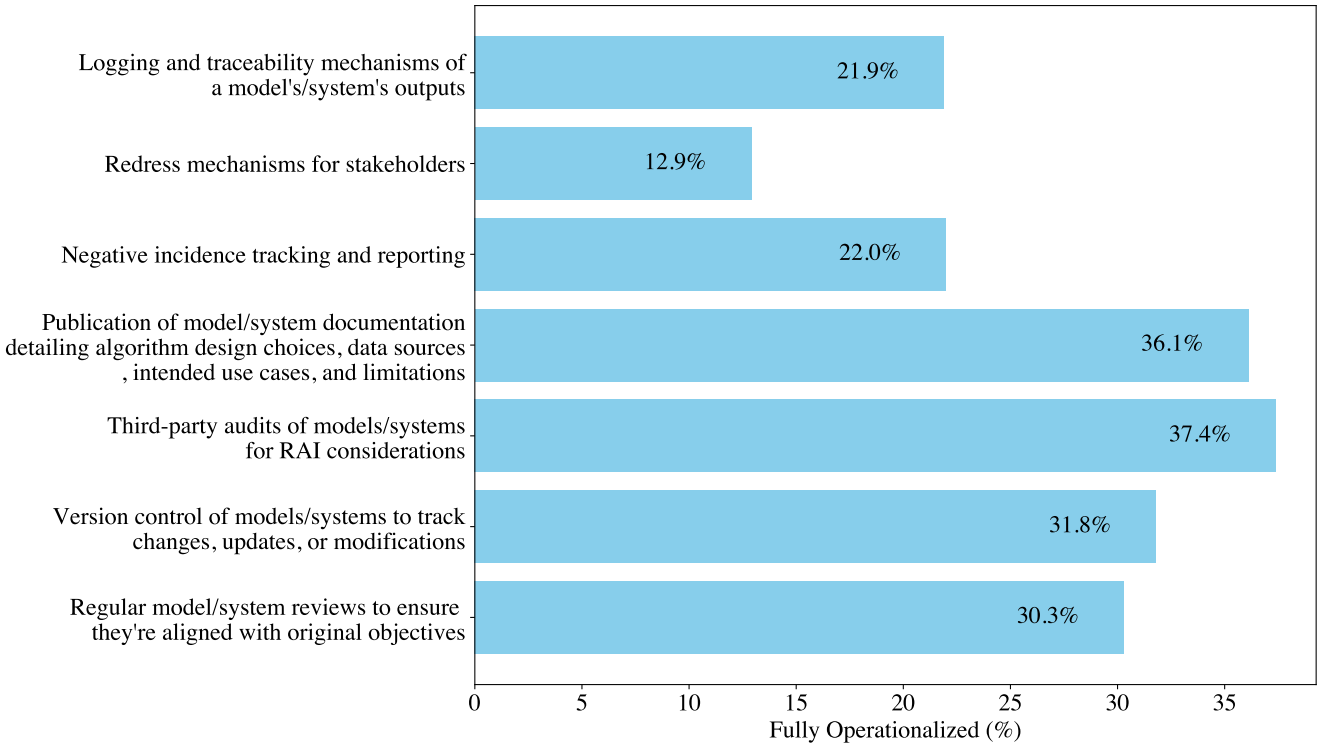

Figure 21: Implemented RAI measures to ensure accountability.

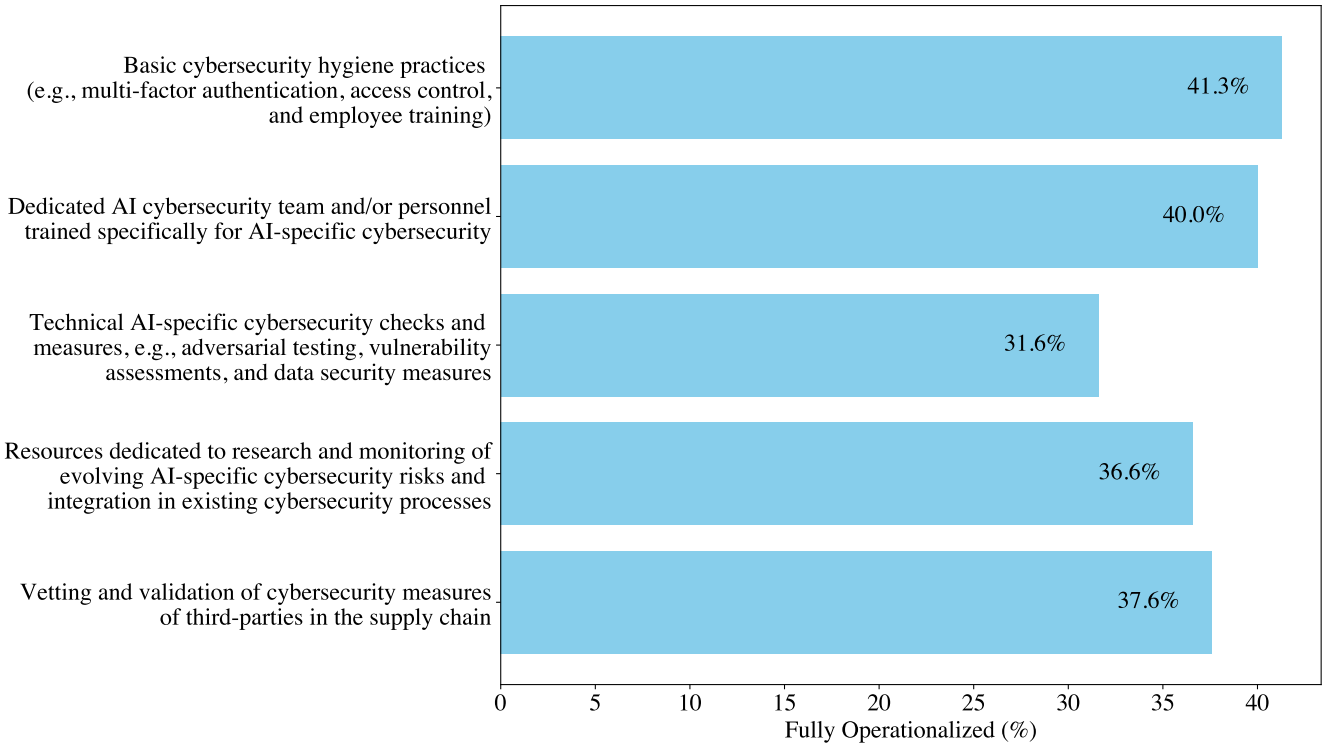

Figure 22: Implemented RAI measures to address cybersecurity risks.

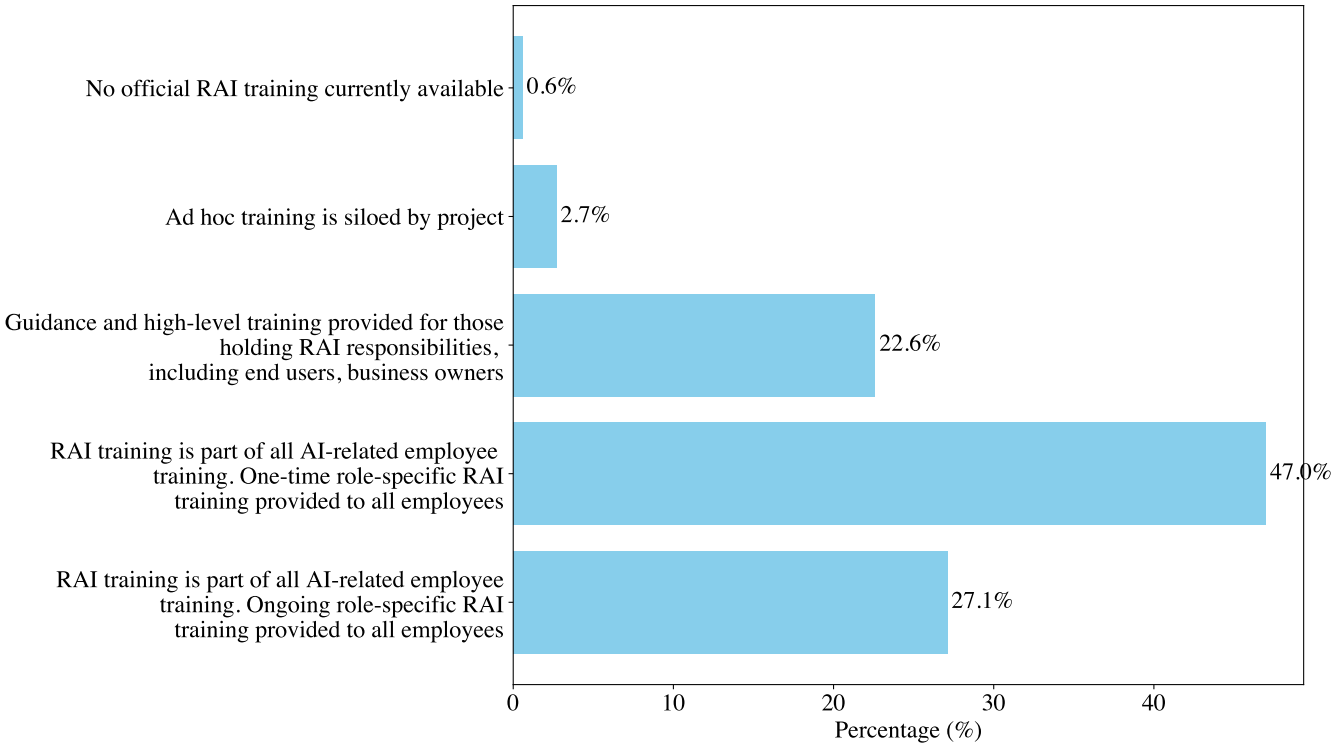

Figure 23: RAI training availability in surveyed organizations.