May I speak? Perceptions on ethical concerns and power while developing software in AI teams

Gabriel M. C. Guimarães
University of Brasília (UnB)
Brasília, Distrito Federal, Brazil
gabrielciriatico@gmail.com

Geraldo Pereira Rocha Filho
State University of Southwest Bahia (UESB)
Vitória da Conquista–BA, Brazil
geraldo.rocha@uesb.edu.br

Gilmar dos Santos Marques
University of Brasília (UnB)
Brasília, Distrito Federal, Brazil
gilmar.marx@gmail.com

Edna Dias Canedo
University of Brasília (UnB)
Brasília–DF, Brazil
ednacanedo@unb.br

ABSTRACT

Context: As Artificial Intelligence (AI) technologies increasingly influence decision-making processes, ethical concerns in AI development have gained significant attention. However, most research in this area focuses on high-level decision-makers or large technology corporations in the Global North, overlooking the lived experiences of software engineers in underrepresented contexts. **Goal:** This study investigates how Brazilian software engineers working in AI teams at a mid-sized company perceive ethical principles, navigate ethical dilemmas, and respond to emerging ethical concerns in their everyday work. The focus is on understanding how Junior, Mid-Level, and Senior developers interpret and apply ethical frameworks during the concrete development of AI systems such as classifiers, image generators, and object detectors. **Method:** A mixed-methods approach was adopted, combining data from 18 survey responses and 8 in-depth semi-structured interviews. The qualitative data were analyzed using a combination of inductive and deductive coding to identify recurring patterns, ethical challenges, and coping strategies. **Results:** Findings indicate that while participants are generally aware of key ethical principles—such as fairness, transparency, and accountability—their ability to act on ethical concerns is limited by factors including organizational hierarchies, lack of formal ethical training, and insufficient autonomy in design decisions. Ethical deliberation is often informal, reactive, and constrained by resource and power asymmetries. **Conclusion:** This research highlights the ethical agency of AI practitioners in the Global South and contributes to expanding the geographical and professional scope of AI ethics literature. It emphasizes the urgent need for organizational structures that empower developers to raise, discuss, and resolve ethical issues throughout the AI development lifecycle.

KEYWORDS

Artificial Intelligence, Ethics, Code of Ethics, Employment, AI Practitioners

1 Introduction

As debates on AI ethics and regulation intensify—particularly around major tech companies like Google, Meta, Microsoft, and OpenAI - attention tends to shift away from the individuals who actually design and build AI systems [14, 19]. Instead, researchers and policymakers often concentrate on corporate representatives, largely overlooking developers' perspectives on their own agency, as well as their understanding and application of ethical principles in practice [4, 11, 17, 24]. Yet, power dynamics in the workplace, the structure of formal education, and the evolving legal landscape significantly influence how software is built—and ultimately, how it affects society [8].

The rise of generative AI models during the COVID-19 pandemic [7, 28], notably with the public release of ChatGPT in November 2022, further accelerated global regulatory efforts. In March 2024, the European Union passed the AI Act [6], establishing categories such as `Prohibited AI Practices`, `High-Risk AI Systems`, and setting obligations for `General-Purpose AI Models` like ChatGPT. In the United States, Executive Order 14110 [21] outlines ethical and safety guidelines for federal agencies deploying AI. Meanwhile, Brazil's Artificial Intelligence Legal Framework (PL 2338/2023) [20], currently under review in the Senate, proposes principles such as non-discrimination, transparency, and respect for fundamental rights, while also classifying AI systems according to their risk levels.

Despite growing concerns about the societal impacts of AI, most regulatory efforts continue to focus on the technology itself, rather than the conditions under which it is developed. Labor-related dimensions of AI production remain largely absent from current policy debates. The EU AI Act [6], for example, offers only limited attention to the working conditions of those involved in designing, testing, or implementing AI systems. It includes a chapter on codes of conduct and guidelines, which encourages ethical practices but stops short of defining or enforcing them. Similarly, Brazil's PL 2338/2023 [20] recommends that organizations develop and promote codes of conduct, yet it does not address labor conditions or workplace ethics in any depth.

The development of AI systems relies on a diverse array of human labor, ranging from highly skilled software engineers to low-paid data annotators [5]. In Brazil, this work often takes place in complex organizational environments, where formal job roles are accompanied by informal expectations and blurred ethical boundaries [2]. Mid-sized companies—frequently overshadowed by global tech giants—play an important role in deploying AI in sectors such as healthcare [12], finance, and logistics. Workers in these settings

may face conflicting pressures: upholding ethical standards, meeting managerial demands, and operating within uncertain legal frameworks. Understanding how AI professionals in Brazil perceive ethics, power, and responsibility in their daily work is thus essential for addressing gaps in both academic research and policymaking.

This study aims to explore how Brazilian software engineers involved in the development of AI solutions understand and apply ethical principles, as well as how they perceive ethical dilemmas in their professional practice within a mid-sized company. We collected data through a survey answered by 18 engineers, focusing on ethical principles and challenges. In addition, we conducted in-depth interviews with 8 of these participants to gain deeper insights into the ethical dilemmas they have encountered throughout their careers. Our investigation is guided by the following research questions (RQ):

RQ.1: **How do AI developers perceive ethical principles?**
RQ.2: **What are AI developers' perceptions on ethical concerns?**
RQ.3: **How do AI developers perceive their own agency when addressing ethical challenges?**

Our main findings show that AI practitioners are primarily concerned with technical ethical principles, such as bias and the protection of sensitive and personal data. Technical challenges also dominate their work, with bias-related errors largely attributed to issues with the data rather than the models themselves. Additional challenges include pressure for rapid delivery and the absence of ethical considerations during the early planning stages. Developers often feel they lack agency to influence broader project decisions—such as timelines and objectives—and are instead confined to implementing technical workarounds. This paper contributes to the discourse on how AI practitioners engage with ethics during software development. It also highlights the importance of ethics education for software engineers. Unlike prior studies that focus on developers in North American big tech companies [18],[22],[13], our work centers on professionals in a mid-sized Brazilian company, where discussions on AI and ethics are still taking shape.

2 Related Work

Several studies have explored how ethical principles are addressed in AI development, with approaches ranging from methodological tools to empirical investigations into practitioners' perceptions. Cerqueira et al. [4] proposed the RE4AI Ethical Guide, a practical tool designed to assist software teams in eliciting ethical requirements in AI-based systems, particularly in agile development contexts. Developed through the Design Science Research methodology, RE4AI comprises 26 cards organized around 11 ethical principles and is implemented as an interactive web-based application. Compared to prior approaches such as ECCOLA, RE4AI provides enhanced support through tool suggestions, explanatory resources, and card filtering features. Empirical results showed that the guide increases ethical awareness and helps translate abstract principles into concrete user stories. This underscores the value of participatory tools that bridge the gap between high-level ethical concepts and day-to-day development activities, especially in environments with low institutional maturity regarding AI ethics.

Widder et al. [26] conducted a mixed-methods study involving 115 survey respondents and 21 interviewees from various sectors and continents. Their work investigates not only the ethical concerns raised by software engineers, but also what actions are taken in response and what factors affect their ability to act. Ethical concerns ranged from technical bugs and feature design to systemic issues, such as business models. Developers responded by proposing technical fixes, negotiating within organizational limits, or refusing to engage. Structural conditions—such as financial precarity, migration status, and organizational culture—were found to significantly shape engineers' capacity to address ethical issues. The study emphasizes that ethical awareness alone is insufficient without the power and resources to act on it.

Borba [3] conducted a systematic literature review to map the ethical principles and guidelines found in both academic and industry discourse on AI ethics. The study identified 26 unique ethical principles, with transparency, fairness, and accountability being the most frequently cited. Transparency was further classified into technical (e.g., model explainability) and operational (e.g., system communication) dimensions. Fairness was often linked to biased data, unequal outcomes, and responsible deployment, while accountability highlighted the importance of defining responsibilities among stakeholders throughout the AI lifecycle.

Khan et al. [11] surveyed 99 participants, including both AI practitioners and policymakers, to investigate their perspectives on ethical principles and challenges. Based on a prior systematic review, the authors validated 15 core ethical challenges, encompassing both technical aspects (e.g., privacy, explainability, data security) and broader socio-ethical concerns (e.g., justice, beneficence, human dignity, non-maleficence). Their study reinforces the multifaceted nature of ethical challenges in AI, as well as the need for alignment between practitioners and regulatory frameworks.

Pant et al. [15] examined the barriers faced by AI practitioners when trying to incorporate ethics into their workflow. Using a mixed-methods survey of 100 participants, the study identified three main categories of challenges: general (e.g., time constraints, cost), technological (e.g., data complexity), and human-related (e.g., cognitive bias, lack of consensus). The findings suggest that ethical awareness is primarily acquired in the workplace (63%), with formal education (43%) and university training (2%) playing secondary roles. Participants reported that implementing privacy and data protection measures is relatively feasible, while adhering to human-centered values was considered particularly difficult (27%).

Vakkuri et al. [23] investigated how ethical considerations are integrated into AI projects through three Finnish case studies: a social marginalization detection tool, an NLP-based diagnostic system, and an indoor navigation application. The study found that ethical concerns were rarely addressed formally. Developers tended to focus on performance and bug-fixing over ethical reflection. Ethical concerns, when raised, often remained personal and un-acted upon, and transparency was frequently deprioritized in favor of system efficiency. None of the projects adopted structured methods for ethical deliberation, revealing a gap between academic discourse and industry practice. Finally, Porto et al. [17] conducted a systematic literature review focusing on techniques for eliciting, analyzing, and specifying ethical requirements in information systems. The review synthesized 47 primary studies, identifying diverse strategies

such as the use of user stories, interviews, ethical modeling, and frameworks like ECCOLA. The review reinforces the importance of integrating ethical reflection into all stages of software development and highlights a growing body of technical approaches aimed at supporting this integration.

3 Study Settings

To address our research questions, we conducted a survey complemented by follow-up semi-structured interviews with software developers working on AI-based systems. This mixed-methods approach allowed us to gather both breadth and depth of insights regarding the practical engagement of practitioners with ethical principles in AI.

The survey instrument, summarized in Table 1, was structured into five sections: (1) Consent and Contact, where participants could voluntarily provide contact information for a potential interview; (2) Demographics, to capture participants' backgrounds and roles; (3) Understanding of AI Ethical Principles, which included a Likert scale assessing familiarity with 26 ethical principles commonly referenced in AI ethics literature; (4) Challenges in Adopting AI Ethical Principles, where participants described ethical dilemmas and obstacles encountered during the development of AI systems; and (5) Perceptions of Ethical Challenges in AI, featuring a Likert scale for ranking 15 specific ethical concerns related to AI practice. All survey questions and answer options are available in Zenodo, including the complete material of this study, at https://zenodo.org/records/17254439.

The ethical principles included in the survey were selected based on the systematic literature review conducted by Borba [3], while the ethical concerns were derived from the SLR by Vakkuri et al. [25]. These two sets of elements formed the basis for two of the core survey sections: Understanding of AI Ethical Principles and Perceptions of Ethical Challenges in AI. The section titled Challenges in Adopting AI Ethical Principles served as a bridge to the qualitative stage of the study. This open-ended question, together with the interview guide, replicates the methodological approach proposed by Widder et al. [26], as our research shares a similar goal: to map developers' perceptions of ethical issues during the development of AI systems. The full interview guide is presented in Table 2.

The only modifications made to Widder et al.'s original guide were the inclusion of questions IQ9, IQ10, and IQ11 (Table 2), which were tailored based on individual survey responses. These questions aimed to explore cases where participants had rated certain principles or concerns at the extremes of the Likert scale—either as having very little or very strong impact on their work—thus enabling a more targeted and reflective qualitative exploration.

The survey and interview guide were initially drafted by the first author and subsequently reviewed by the second author, a researcher with over 20 years of experience in conducting empirical studies. Based on her feedback, the instruments were revised and finalized for deployment. Prior to full distribution, a pilot study was conducted with three professionals to assess the clarity and usability of the materials; these responses were excluded from the final analysis.

The survey was made available between November 1st, 2024, and April 1st, 2025, and was conducted within a mid-sized Brazilian technology company (100–1000 employees) actively engaged in AI-related projects. The company granted formal authorization to conduct the research, and the survey was distributed remotely via a secure link. All participants who consented to further contact during the survey were invited to participate in follow-up interviews, which were conducted between January 7th and April 4th, 2025. This study received ethical approval from the Ethics Committee of the Institute of Human and Social Sciences at the University of Brasília (approval number 7.575.474). Informed consent was obtained from all individual participants as well as from the participating organization.

Demographic information is partially summarized in Table 3. A total of 18 developers completed the survey, and 8 of them participated in follow-up interviews. All participants were Brazilian. The distribution by state (SQ4 of Table 1) was: São Paulo (27.8%, 5), Pernambuco (16.7%, 3), Distrito Federal, Ceará, and Bahia (11.1% each, 2), and Piauí, Alagoas, Minas Gerais, and Paraíba (5.56% each, 1). Age groups (SQ5) were evenly split: 50% (9) were aged 18–24, and 50% (9) were aged 25–34. Regarding race (SQ7), 72.2% (13) identified as White, 22.2% (4) as Brown, and 5.56% (1) as Black. In terms of gender (SQ6), 88.9% (16) were men and 11.1% (2) were women.

As for educational background (SQ8), 72,3% (13) had completed an undergraduate degree, 16.7% (3) a master's degree, 5.56% (1) had completed a master's degree, and 5.56% (1) was pursuing a Ph.D. In terms of job seniority (SQ9), 44.4% (8) were mid-level, 27.8% (5) junior-level, 16.7% (3) senior-level, and 11.1% (2) interns. Software industry experience was distributed as follows: 44.4% (8) had 1–3 years, 33.3% (6) had 4–6 years, 16.7% (3) had less than 1 year, and 5.56% (1) had 13–15 years. AI-specific experience (SQ10) was 33.3% (6) with 1–3 years, 33.3% (6) with 4–6 years, 16.7% (3) with 7–9 years, and 16.7% (3) with less than 1 year.

The survey's closed-ended questions were analyzed using descriptive statistics. Open-ended responses were present only in the Challenges in Adopting AI Ethical Principles section and were used to guide the formulation of interview questions for participants who agreed to be interviewed. Interviews were recorded, transcribed, and coded following the methodology proposed by Pant et al. [15] and Vakkuri et al. [25]: initially, inductive codes were generated from the transcripts, which were later grouped into deductive codes through iterative reading. Additionally, interviews produced Key Findings (KFs), as defined by Pant et al. [15], which represent relevant insights and recurring discussion points that emerged from the interviews.

4 Results

The research questions were addressed through data obtained from survey and interviews. Table 4 presents the inductive and deductive codes derived from the interview transcripts. The survey were analyzed using descriptive statistics.

4.1 RQ1: What are AI developers' perceptions on ethical principles?

Among the ethical principles that developers reported they strongly agree they are familiar with or apply in practice, the most frequently

Table 1: Survey Questions

Section	Question
Consent and Contact	SQ1. Do you consent to participate in this research? [Yes/No]
	SQ2. Are you available to participate in a 30-minute interview to further discuss the answers given in this questionnaire? [Yes/No]
	SQ3. If yes, please leave your email or phone number for contact. [Open]
Demographics	SQ4. In which Brazilian state (Federative Unit) do you currently live?
	SQ5. What is your age group?
	SQ6. What is your gender identity?
	SQ7. What is your race/ethnicity?
	SQ8. What is your highest level of education?
	SQ9. How many years of experience do you have in the software industry?
	SQ10. How many years of experience do you have in AI?
	SQ11. What is your current area of expertise?
	SQ12. What is your current job position?
	SQ13. What is the nature of your work/organization?
Understanding of AI Ethical Principles	SQ14. Please indicate how much you agree with the following statement: "Regarding ethical principles in AI, I am familiar with and/or apply these principles in my daily software development activities." [Strongly disagree/Disagree/Neutral/Agree/Strongly agree] (1) Accuracy, (2) Autonomy, (3) Social well-being, (4) Beneficence, (5) Trustworthiness, (6) Human dignity, (7) Diversity, (8) Effectiveness, (9) Fairness, (10) Explainability, (11) Interpretability, (12) Justice, (13) Legality, (14) Non-discrimination, (15) Non-maleficence, (16) Predictability, (17) Privacy, (18) Prosperity, (19) Accountability, (20) Robustness, (21) Security, (22) Data governance, (23) Solidarity, (24) Human oversight, (25) Sustainability, (26) Transparency
Challenges in Adopting AI Ethical Principles	SQ15. Have you ever had ethical concerns when developing AI-related software?[Yes/No]
	SQ16. If yes, what were you asked to do in that project?
	SQ17. If you had ethical concerns, how were they resolved in practice?
	SQ18. Could you list the ethical concerns you had?
	SQ19. What actions did you take as a result of these concerns?
	SQ20. How did your organization or software development team handle your ethical concerns?
	SQ21. How did you feel about the outcome of your actions?
	SQ22. Can you describe anything that made it easier to address your concerns?
	SQ23. Can you describe anything that made it harder to address your concerns?
Perceptions of Ethical Challenges in AI	SQ24. To what extent do the following challenges hinder the implementation of ethical principles in AI? [Very Little/Little/Not at all/Reasonably/Very much] (1) Lack of knowledge on ethics; (2) Vague principles; (3) Generic principles; (4) Practical conflicts; (5) Different interpretations; (6) Lack of technical knowledge; (7) Structural limitations; (8) Lack of legal monitoring bodies; (9) Lack of legal framework; (10) Business interests; (11) Plurality of ethical methods; (12) Ethical dilemmas; (13) Machine distortion; (14) Lack of guidelines; (15) Lack of multicultural cooperation.

cited were trustworthiness (67.7%), privacy (61.1%), and accuracy, non-maleficence, effectiveness, explainability, interpretability, and security (55.6% each). On the other hand, the principles with the lowest rates of strong agreement were prosperity (22.2%), sustainability and social well-being (33.3% each), and human oversight, autonomy, and accountability (38.9% each).

A detailed view of respondents answers can be seen in Figure 1. This chart shows participants' responses on a 5-point Likert scale regarding their level of agreement with the statement: "I am familiar with and/or apply this principle in my daily software development activities." The scale ranges from Strongly Disagree to Strongly Agree. The figure displays the percentage distribution of responses for each AI ethical principle, such as Transparency, Fairness, Accountability, and others.

These results indicate a stronger familiarity and application among developers of principles associated with technical solutions, such as trustworthiness, which was defined in the survey as trust in the system's ability to operate consistently, and accuracy, understood as a commitment to minimizing system errors. Meanwhile, principles more closely linked to social and non-technical aspects, such as prosperity (promotion of economic development and social well-being), sustainability (commitment to the environment), and social well-being (promotion of conditions that benefit society at large), were less frequently recognized. Notably, sustainability was the only principle where a respondent strongly disagreed with its application.

Concerns about bias. Some developers raised significant concerns about biases emerging from AI systems, particularly in relation to how socioeconomic profiles were portrayed. One developer

Table 2: Interview Guide Questions

Id	Question	Notes
IQ1	To start, can you tell me a little about yourself and your professional background?	General opening
IQ2	Let's talk about the experience you mentioned in the form regarding facing ethical challenges. You said you were asked to [summary of task]. To begin, can you tell me what your employment type was in that project (e.g., full-time, contractor, researcher)?	If answered "yes" to ethical dilemmas in SQ15
IQ3	In the form, you wrote that you were concerned with [summary of concerns]. Can you tell me how these concerns emerged?	Same as above
IQ4	Do you think there was anything you could do, within your responsibilities, to address these ethical issues? Or was it beyond your scope?	Same as above
IQ5	Looking back, why did you perceive this as an ethical concern? How did you come to see it that way?	Same as above
IQ6	You mentioned [facilitating factor] helped you act in the project. Can you tell me more about that? Was there anything else that helped?	Same as above
IQ7	You also mentioned [hindering factor] made it more difficult to act. Can you elaborate? Was there anything else that made it harder?	Same as above
IQ8	Do you know any colleague who has gone through something similar?	Same as above
IQ9	Why did you say that [insert principles] are ones you disagree you know or apply in your daily AI development activities?	Follow-up from survey
IQ10	Why did you say that [insert principles] are ones you neither agree nor disagree you know or apply?	Follow-up from survey
IQ11	Why did you say that [insert challenges] are ones that significantly hinder the implementation of ethical principles in AI?	Follow-up from survey

Table 3: Demographic and professional profile of participants

Age Group	Gender	Race	Degree	Years in Software	Years in AI	Role	Interviewed	Practical Concerns
18–24	Men	White	Graduated	1–3	1–3	Junior Machine Learning Developer	Yes	Privacy
18–24	Men	White	Graduated	1–3	4–6	Senior Machine Learning Tech Lead	Yes	Bias, efficiency
25–34	Men	White	Master Student	<1	4–6	Mid-level Machine Learning Developer	No	
25–34	Men	White	Master	4–6	4–6	Mid-level Machine Learning Developer	Yes	Intellectual property, hierarchy
25–34	Men	White	Master Student	4–6	7–9	Senior Project Manager	Yes	Privacy, legal
25–34	Men	White	Graduated	1–3	4–6	Mid-Level Machine Learning Tech Lead	Yes	Efficiency
25–34	Men	Brown	Graduated	13–15	7–9	Mid-level Machine Learning Developer	No	
18–24	Men	White	Graduated	1–3	1–3	Junior Machine Learning Developer	Yes	Privacy
18–24	Men	Black	Graduated	<1	<1	Machine Learning Intern	No	
18–24	Men	White	Graduated	1–3	1–3	Junior Machine Learning Developer	No	
18–24	Women	Brown	Graduated	<1	1–3	Junior Machine Learning Developer	No	
18–24	Women	White	Graduated	1–3	<1	Mid-level Front-End Developer	No	
25–34	Men	White	Master Student	4–6	4–6	Mid-level Machine Learning Developer	Yes	Bias, prejudice
18–24	Men	Brown	Graduated	1–3	<1	Junior QA/Testing Intern	No	
25–34	Women	White	Graduated	4–6	1–3	Mid-level Back-End Developer	No	
25–34	Men	White	Graduated	1–3	1–3	Junior Back-End Developer	No	
18–24	Men	White	Graduated	4–6	4–6	Junior Machine Learning Developer	No	
25–34	Men	Brown	PhD Student	4–6	7–9	Senior Business Architecture	Yes	Bias, prejudice, intellectual property

explained, "The idea behind the project was to generate images for personas, profiles of people, from high-income and low-income backgrounds, from urban, metropolitan, and peripheral areas. And then the AI was responsible for the generation, right? In some cases, for profiles of low-income individuals from the periphery, the AI would often depict the person as Black. This is a prejudice embedded within the model—assuming that if someone is low-income and from the periphery, they must be Black. This happened repeatedly during testing, and it was a reality shock for us when dealing with these models" (Interviewee 3).

Such accounts reveal a critical awareness among developers about the replication and amplification of social biases through AI outputs. The realization that biases could unintentionally arise from the training data itself underscores developers' ethical concerns. It also points to the challenges they face in designing AI systems that do not perpetuate harmful stereotypes, particularly in projects dealing with sensitive or representational outputs like user profile images.

Concerns about sensitive and personal data. Handling sensitive and personal data emerged as a major concern among developers, particularly in contexts where data quality required manual intervention. One developer shared, "Some documents weren't in a good enough format to feed directly into the model. You had to adjust them manually—correct orientation, among other things. Automated models couldn't handle all cases, and that's when I realized I was dealing with a lot of personal, sensitive content—various types of documents, passports, IDs, and so on. It made me think, 'Isn't there a safer way to handle this?'" (Interviewee 1).

Other developers emphasized the gravity of working with sensitive information, noting the additional layers of responsibility involved. "We work with personal and sensitive data. Personal information is already problematic because you can identify individuals. But with sensitive data, the stakes are even higher—you deal with health patterns, social preferences, purchasing habits. The misuse of such data can lead to disastrous consequences" (Interviewee 6). Another participant recalled a project where an API was used to

Table 4: Deductive and Inductive Codes

Deductive Code	Inductive Codes	#
Focus on results	Priority for efficiency, Short development time, Practice over ethics, Client focus	16
Data concerns	Concerns with sensitive data, Questions about user consent, Concerns with personal data	13
Technical solutions for ethical dilemmas	Technical solution, Lack of technical knowledge, Explainability for accountability, Insecurity due to technical issues	11
Concerns about model bias/prejudice	Concern with bias, Concern with prejudice, The machine is not guilty, the developer or data is	10
Ethical concerns before execution	Legal responsibility for safety, Concerns before execution, Ethical guidance for project, Ethical questioning as habit	8
Ways to ease ethical anxiety	Responsibility shifted to companies, Tradition dulls concern, Secondary data is not personal data, Developer good faith, Outsourcing concerns	7
Lack of space to discuss ethics	Isolation, Developer unaware of project details	6
Existence of ethical discussion spaces	Intercultural diversity prompting questions, Openness to ethical discussion, Multidisciplinary ethics network, Dialogue with team	6
Unavoidable ethical problems	Concern during execution, Ethical concern due to business needs	6
Generative AI/LLMs specific issues	Intellectual property concern, Difficulty solving LLM issues, LLM-specific problems	6
Lack or unawareness of AI laws	Lack of AI legislation knowledge, Big players as guides, Lack of regulations	5
Developer's hierarchical position	Developer submission to hierarchy, Lack of voice	5
Issues from outsourced developers	Developer company more concerned than client, Outsourced development	5
Ad-hoc solutions for ethical issues	Personal temporary solution	4
Positive impact of legislation	Ethical guarantees for legal security, LGPD impact	4
Ethics seen as unambiguous	Ethics as universal values	3
Areas less prone to ethical dilemmas	Existence of areas with fewer ethical concerns, Academia seen as less problematic	3
Avoidance of questioning project purpose	Inability to cancel project, Social well-being beyond scope	3
Ethics seen as ambiguous	Ambiguity around ethics	2
Beyond technical solutions for ethics	Communication as solution	2
Concerns about AI project impact	Concern with practical model impact	1
Use of methodologies to address ethical concerns	Use of AI development methodologies	1
Negative or null impact of legislation	Legislation has no concrete impact	1
Unable to quit project due to career concerns	Lack of perspective in other companies	1

collect detailed personal information through users' CPF numbers: "The API returned full names, mothers' names, cities of birth—lots of highly personal data. We had to be extremely cautious about potential leaks because we had signed confidentiality agreements. Any leakage would have made us personally liable" (Interviewee 4).

RQ.1 Summary: AI developers are more familiar with and tend to apply ethical principles related to technical dimensions, such as trustworthiness, privacy, and accuracy. In contrast, principles tied to broader societal or non-technical concerns—like prosperity, sustainability, and social well-being—were less recognized. This suggests a stronger alignment with operational and performance-oriented values, while social and environmental considerations remain underprioritized in everyday practice.

4.2 RQ.2: What are AI developers' perceptions on ethical concerns?

When asked about the challenges that hinder the implementation of ethical principles in AI development, developers most frequently pointed to business interests (61.1%) as the most impactful, followed by practical conflicts and the absence of legal monitoring bodies (both at 44%). In contrast, machine distortion (16.7%) was among the least cited concerns, along with factors such as lack of knowledge on ethics, limited multicultural cooperation, absence of a legal framework, plurality of ethical methods, ethical dilemmas, and again, lack of legal monitoring (5.6% each). A detailed view of respondents answers can be seen in Figure 2. A detailed view of respondents' answers is presented in Figure 2, which shows the distribution of Likert scale responses (ranging from 'Not at all' to 'Very much') for each ethical challenge identified.

Practical challenges related to business needs and organizational hierarchies were particularly prominent. Business interests were defined in the survey as reconciling profit motives, delivery speed, and other business priorities with ethical standards. Practical conflicts

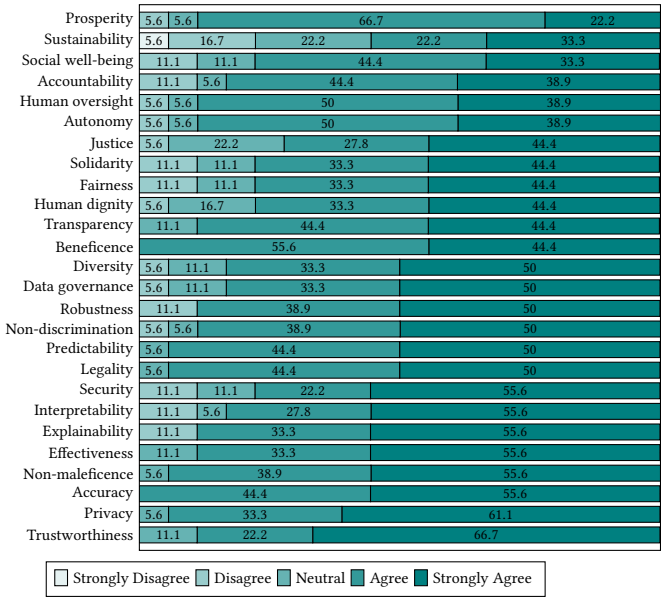

Figure 1: Distribution of participants' responses regarding their familiarity with and application of AI ethical principles in daily software development.

referred to clashes between ethical principles and project priorities. More theoretical concerns about ethics were not seen as a major barrier: while 94.5% of developers agreed that generic principles impact implementation at least reasonably, 27.8% stated that lack of knowledge on ethics has little or very little impact.

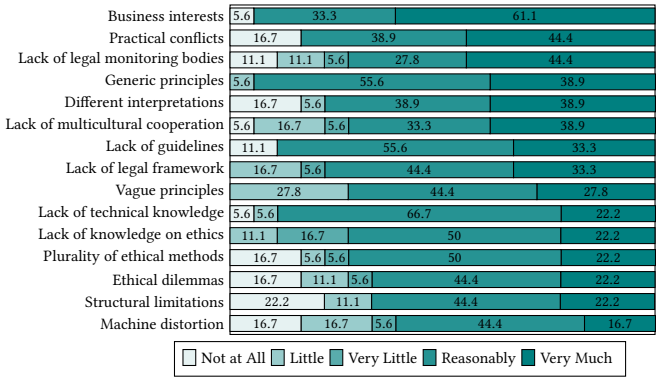

Figure 2: Likert scale responses for each AI ethical challenge.

Machine as an innocent actor. Machine distortion—defined as distortions arising throughout the pipeline, from input data to model outputs—was perceived by many developers as having minimal impact. In total, 39% indicated it affects ethical implementation 'not at all', 'little', or 'very little'. This perspective reflects a widespread belief that machines merely reproduce patterns found in the training data, rather than introduce bias independently. As one participant noted, "Machine distortion has little impact. Everything AI outputs comes from the datasets we choose. For example, if we train AI on a social network with a lot of bias, obviously the AI will exhibit bias—but that's not the AI's fault; it's about how we selected the data" (Interviewee 8).

This perspective resonates with how developers perceive model bias more broadly: the problem is seen as originating from the data, not the algorithms themselves. As another developer put it, "Today we haven't faced it directly, but it's something we're concerned about—having a good balance of input data to avoid biased decisions" (Interviewee 6). No interviewee suggested that algorithms could independently cause bias, such as through structural reinforcement of certain behaviors regardless of data.

Priority on deliverables. Pressure to achieve quick results and deliver projects on time was repeatedly mentioned as a barrier to the implementation of ethical principles. According to one developer, "Ethical concerns are very complicated to address in a project because we are always working for a client who wants quick results. They are not concerned about ethics; they are concerned about the impact on their day-to-day business" (Interviewee 2).

Raising ethical concerns during a project can even be perceived as introducing additional costs, both in terms of time and quality. One developer mentioned the difficulties of proposing additional efforts to ensure data anonymization: "We had to make an unplanned effort to address the personal and sensitive data issues the client gave us. If you tell the client and the managers that you need an extra week in a one-month project just to deal with anonymization, sometimes they'd rather accept the risk, no matter how serious it might be" (Interviewee 6).

Ethical concerns arising during development. Ethical concerns reported by developers emerged at different stages of development, either during the planning phase or during execution. While three developers reported that concerns arose when project results started indicating a problem, another three commented that they habitually map out potential problematic aspects during the planning stage. Among the first group, ethical issues such as bias and prejudice became concerns when they negatively affected the client's expected outcomes: "For example, let's say it was a study for a tech company, so they were looking at purchase profiles related to that company. And then they created tactics based on representative people from these groups. Everything was going fine, but we entered a very specific market, because they work in the gaming market, right? And then, in all the research, all the clustering, it always picked up a predominantly male profile, right? So, it was a very strong bias that we faced in that project" (Interviewee 2).

This developer started noticing biased results, focusing predominantly on male profiles, while the client was concerned with reaching individuals of other genders as well. Among those who reported concerns from the planning phase, issues were mainly related to data handling. "No, we already knew what kind of data we were going to work with, right? We knew there would be sensitive information, so it was something we tried to address in advance, but our recommendations were not heard." (Interviewee 6) "I think the concerns arose during the model training process, right, especially in the labeling phase. There is a very manual part, where you are dealing with the data and have to label each piece of information. So, I think it was during that manual process that I realized some documents were not in a good enough state to input directly into the model, so I had to fix them—like adjusting from horizontal to

vertical, and making other very manual adjustments " (Interviewee 1).

Ethics as universal or ambiguous concepts? Developers expressed differing views on whether ethical notions are universal or inherently ambiguous. On the one hand, one developer, commenting on a project he refused to participate in out of fear that the facial recognition model would reproduce racial biases against Black individuals, remarked: "I think that, for example, when it comes to putting ethics into practice, it is something — if I'm saying something wrong, correct me — but I think it's very new. So it's very complicated to say what is ethical for one person — would it be for another? Sometimes, for one person, implementing a facial recognition algorithm would not be a problem because they don't see it that way" (Interviewee 5).

On the other hand, two developers stated that they see ethical ambiguity as having little impact on the implementation of ethical principles, suggesting there is broad consensus among developers about what is right and wrong. One explained: "I think I mentioned that it has very little impact because everyone has a notion of what is ethical or not, right? We know what is right and what is wrong, especially when you work in the area. You have that notion of right and wrong, but whether you apply it or not is another question. So I think it's kind of general knowledge. Maybe if you're unsure about a tool, you can ask someone, but you still know what is right or wrong" (Interviewee 1). Another added: "For example, I think vague principles make it only slightly harder because everyone knows what ethical principles they need to have in their modeling today, right? Especially when it comes to generative agents, I think everyone knows they must not be prejudiced, must not be aggressive, must not carry data biases. I think vague principles are already very well understood, except perhaps in more open-ended research scenarios. And I also think that different interpretations of these principles have little impact" (Interviewee 2).

> **RQ.2 Summary**: Developers perceive commercial interests and practical conflicts as the main barriers to implementing ethical principles in AI. Issues like machine bias and lack of ethics knowledge are seen as less impactful. Many attribute bias to input data rather than algorithms. Time pressure and fast delivery demands hinder ethical considerations. While some see ethics as ambiguous, others believe there is general agreement on what is right or wrong.

4.3 RQ.3: How do AI developers perceive their own agency when addressing ethical challenges?

Eight out of eighteen developers (44.4%) reported having faced an ethical dilemma during the development of an AI project. On average, these developers had more experience with AI (5.57 years) than those who had not encountered such dilemmas (4.1 years). This trend also holds for general software development experience, as shown in Figure 3. Developers earlier in their careers reported facing ethical dilemmas less frequently than their more experienced peers. In fact, developers with less than one year of experience were more than twice as likely to say they had never encountered an ethical dilemma compared to those who had.

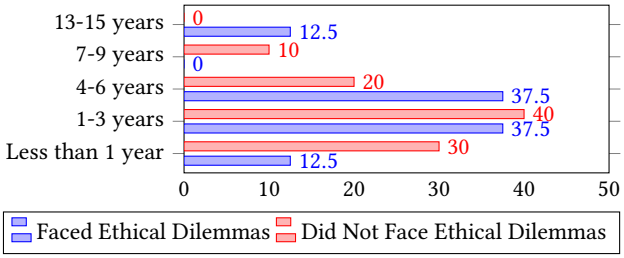

Figure 3: Comparison of years of software development experience between developers who have and have not faced ethical dilemmas.

Spaces for discussion. Only one developer mentioned having a dedicated space within their team to discuss ethical issues related to the project. This project, which involved automating the completion of police reports, was led by a coordinator who emphasized the sensitivity of the data and the importance of mitigating risks: "Over time, our coordinating professor gave attention to these issues, really emphasizing that we were dealing with sensitive data, data that wasn't ours. We had to take actions to avoid impacts, like using a test CPF with fictitious data for testing purposes" (Interviewee 4).

Another developer reported discussing ethical concerns with friends outside of work, while a third mentioned working alone on the data science component, which made it difficult to find colleagues who could fully understand the ethical implications: "The team was very lean. I exposed these concerns to my colleagues, but I was mainly the one handling the data science part" (Interviewee 6). In the absence of an open and prepared environment for such discussions, developers often chose not to share their concerns and instead sought individual solutions without collective support.

Hierarchical conflicts were identified as barriers both to resolving practical dilemmas and to questioning broader ethical aspects of projects, such as social impact or beneficence. Developers frequently emphasized their subordinate role in the organizational hierarchy. One interviewee explained: "Usually the developer is under a hierarchy where they must develop regardless of what they think. That's how it works; they will develop what the company or project demands. And if they express discomfort, what protects them? What ethical or legal principle supports an employee who says 'This is completely wrong, I won't do it'? There's nothing protecting them today; you have to do it" (Interviewee 3).

Another developer echoed this sentiment, noting that although objecting to a project might be theoretically possible, technical team members usually lack the authority to stop or pause a project - especially when financial concerns are involved: "At the time, I was just a developer. It's not that it couldn't have happened, but on the technical side of the team, protesting… It's exactly that — not having the decision-making power to stop or pause the project, because it involves financial matters and everything else" (Interviewee 6).

When reflecting on a project involving artistic style transfer for user-submitted images, one developer expressed frustration at their

inability to question problematic aspects of the work: “Even if I had a voice, it wouldn’t have changed anything, because it was something that came from higher up — I don’t even know if it’s called stakeholders — but they mapped internal themes and research priorities. This project was a priority for them, so no matter what a junior researcher said, it wouldn’t change anything” (Interviewee 8).

RQ.3 Summary: AI developers who reported facing ethical dilemmas tended to have more experience than those who had not. Less experienced developers were less likely to recognize or report such dilemmas. Most participants lacked structured spaces within their teams to discuss ethical concerns. Hierarchical constraints often limited developers’ ability to question or challenge ethically problematic decisions. As a result, many developers navigated ethical issues individually and without institutional support.

5 Discussion

Prioritizing technical principles: convergence with literature. The results of this research significantly align with previous studies regarding developers’ preferences for technical ethical principles. Khan et al. [11] identified transparency, accountability, and privacy as the most commonly recognized principles among AI professionals, converging with our findings that trustworthiness (67.7%), privacy (61.1%), and accuracy (55.6%) were the most widely applied by participants.

This convergence suggests a consistent pattern in developers’ perceptions: there is greater familiarity and comfort with principles that can be translated into measurable technical solutions. As observed in our results, principles such as prosperity (22.2%) and sustainability (33.3%) - which require broader social considerations - received lower adherence, echoing Agbese et al. [1] findings that ethical components related to societal and environmental well-being are perceived as less profitable and therefore less prioritized.

Practical barriers vs. knowledge: confirming established patterns. Our findings regarding the main barriers to ethical implementation - particularly business interests (61.1%) and practical conflicts (44%) - find strong support in the literature. Khan et al. [11] identified "conflict in practice" as the most severe challenge for AI ethical principles, while Agbese et al. [1] highlighted that the financial value of ethical requirements is “practically non-existent.”

Particularly revealing is our finding that only 27.8% of developers consider lack of ethical knowledge as a significant barrier. This contrasts with the common perception that developers lack ethical training and suggests, as argued by Agbese et al. [1], that executives and professionals are aware of ethical requirements but face structural pressures that prevent their implementation.

One participant’s narrative - “Ethical concerns are very complicated to address in a project because we are always working for a client who wants quick results” (Interviewee 2) - directly illustrates previous studies’ observations about how temporal and financial pressures override ethical considerations in professional practice.

Outsourcing and responsibility diffusion: expanding the literature. The outsourced nature of the studied company revealed a dimension little explored in the literature on ethical agency in AI. Our results show how outsourcing arrangements create additional layers of responsibility diffusion, with developers feeling even more distant from strategic decisions.

This finding expands observations from previous work on ethical responsibility. While Pant et al. [16] mention uncertainties about responsibility, our findings specify how organizational outsourcing structures intensify this ambiguity. One developer noted: "it becomes unclear to developers where responsibility ultimately lies" - a concern that goes beyond traditional internal hierarchies to include multiple external stakeholders.

This dialogues with Agbese et al. [1] observations about how “multi-stakeholder participation” can result in inadequate governance and enable “ethics washing,” but our research reveals specific mechanisms through which this occurs in outsourcing contexts.

Machines as innocent actors: reinforcing technological neutrality. Our findings regarding machine distortion being perceived as having minimal impact (39% indicated low or no impact) align directly with literature on technological neutrality. As Griffin et al. [10] reported, many professionals adhere to the value neutrality thesis, viewing technology as “neither good nor bad on its own.”

One participant’s perspective - “Everything AI outputs comes from the datasets we choose [...] that’s not the AI’s fault; it’s about how we selected the data” (Interviewee 8) - perfectly exemplifies this algorithmic neutrality view identified in the literature. This perception has important implications for how developers conceptualize their own ethical responsibility: by externalizing the source of bias to data, they may feel less responsible for the ethical consequences of the systems they develop.

Limited agency vs. veiled agency: hierarchical nuances. Our results regarding hierarchical limitations contrast interestingly with some findings in the literature. While Griffin et al. [10] found that more developers (n=17) felt they had authority to intervene than not (n=7), our research reveals a more restricted perception of agency, particularly among junior developers.

However, our findings about “technical solutions as the main pathway” align with the concept of “veiled agency” identified in the literature - where developers exercise ethical agency through technical choices without necessarily explicitly recognizing the ethical dimensions of these decisions. One participant described: “We created a kind of separator there, to cluster data both for female and male genders separately” (Interviewee 2) - a technical solution that addresses ethical concerns indirectly.

Technical solutions: preference confirmed by literature. The strong preference for technical solutions (88.9% consider lack of technical knowledge as the greatest obstacle) finds direct parallel in the findings of Pant et al. [16], where developers consistently prioritized technical interventions for ethical issues. Our participants frequently resorted to preprocessing, prompt engineering, and anonymization as primary responses to ethical dilemmas.

This tendency may reflect both developers’ technical training and a coping strategy that allows continuing projects without directly confronting more fundamental ethical questions. As observed in the literature on Pant et al. [16], developers may prefer technical solutions because these are perceived as more controllable and measurable than broader social or organizational interventions.

Absence of discussion spaces: converging with identified patterns. Our finding that only one developer reported dedicated spaces for ethical discussion contrasts with some literature results. Griffin et al. [9] found that most participants (n=15) would seek colleagues for help with ethical dilemmas. This discrepancy may reflect cultural, organizational, or seniority differences between the studied samples.

Our participants frequently reported navigating ethical issues individually, which aligns with literature observations about junior developers having lesser ethical agency. The lack of structured spaces for ethical discussion may perpetuate dependence on individual technical solutions, limiting opportunities to address ethical issues more systemically.

Implications for governance and regulation. Our findings regarding demand for regulation (88.9% agree that lack of ethical guidelines impedes implementation) converge with Khan et al. [11], who identified "no legal frameworks" and "lacking monitoring bodies" as high-ranking barriers. One participant's perspective on LGPD - "A strong defense would be to show our efforts to anonymize the data" (Interviewee 6) - illustrates how regulations can function not only as restrictions but as support for ethical arguments in organizational contexts.

This observation suggests that regulations can empower developers by providing external justifications for ethical practices, addressing some of the hierarchical and commercial pressures identified both in our study and in existing literature.

6 Threats to Validity

We discuss threats to the validity of our study following the classifications proposed by Wohlin et al. [27]. **Conclusion validity** refers to the extent to which the conclusions drawn from the data are credible. The small sample size (18 survey participants, of whom only 8 agreed to be interviewed) may limit the robustness of our findings and increases the risk of random variation influencing the results. While we performed careful qualitative analysis and triangulated survey and interview data, the limited number of participants restricts the statistical power and generalizability of some findings. **Internal validity** concerns causal relationships. Since our study is exploratory and descriptive, it does not aim to establish causality. However, the potential influence of confounding factors (e.g., participants' prior exposure to ethical training or specific organizational cultures) could affect how developers perceive and report ethical dilemmas.

Construct validity addresses whether the study measures what it intends to measure. To mitigate this threat, we designed our survey and interview protocols based on existing literature on AI ethics and developer agency. However, participants may have interpreted key terms (e.g., "ethical dilemma" or "agency") differently, which could influence their responses. We attempted to reduce ambiguity by providing brief explanations during interviews when needed. **External validity** relates to the generalizability of our findings. Given the relatively small and self-selected sample, caution should be exercised when extending the results to broader populations of AI developers. Participants may have had a particular interest in ethical issues, which could introduce self-selection bias. Future studies with larger and more diverse samples are needed to validate and expand upon our findings.

7 Conclusion

In this study, we examined how Brazilian AI developers perceive ethical principles, concerns, and their own agency in dealing with ethical challenges, through survey data and in-depth interviews. Regarding RQ.1, developers primarily associated ethics with technical principles such as trustworthiness, accuracy, and privacy—reflecting a view that prioritizes performance and compliance over broader societal values.

For RQ.2, participants reported struggling to apply ethical principles in practice, often due to organizational barriers such as rigid hierarchies, lack of formal spaces for ethical reflection, and limited influence over project goals. RQ.3 revealed that developers' perceived agency was shaped by experience, organizational context, and team dynamics. While some found ways to resist or mitigate ethically problematic work—by seeking peer support, proposing technical solutions, or refusing tasks—many felt constrained by their subordinate roles and the absence of institutional support.

Our findings also raise important directions for future research. These include the need to explore accountability gaps in outsourced AI development, where ethical responsibility is often diffuse, and the inadequacy of purely technical approaches when ethical tensions involve the core objectives of AI systems, such as fairness or social impact.

ARTIFACT AVAILABILITY

The supporting data for this work is available at https://zenodo.org/records/17254439.

ACKNOWLEDGMENTS

This study was financed in part by the Project No. 514/2023 – Call Nº 10/2023 – FAPDF Learning Program, a strategic development initiative in the macro areas of Agro Learning, Bio Learning, Gov Learning, and Tech Learning; and Conselho Nacional de Desenvolvimento Científico e Tecnológico CNPq (Grant Nº 300883/2025-0).

Declaration of generative AI in scientific writing

We employed generative AI in the writing process, especially aiming to improve the text quality. We asked for help to improve the text readability and to find suitable synonyms for words based on the context.

REFERENCES

[1] Mamia Agbese, Rahul Mohanani, Arif Ali Khan, and Pekka Abrahamsson. 2023. Implementing AI Ethics: Making Sense of the Ethical Requirements. arXiv:2306.06749 [cs.CY] https://arxiv.org/abs/2306.06749

[2] Kenzo Soares Seto and. 2025. AI From the South: artificial intelligence in Latin America through the sociotechnical imaginaries of Brazilian tech workers. *Globalizations* 0, 0 (2025), 1–16. https://doi.org/10.1080/14747731.2025.2465166

[3] João Gabriel Rossi de BORBA. 2025. *Do Abstrato ao Concreto: Uma Ferramenta para Implementação de Requisitos Éticos em Inteligência Artificial através de Histórias de Usuário.* Dissertação (Mestrado em Informática). Universidade de Brasília, Brasília. https://repositorio.unb.br/handle/10482/51946

[4] José Antonio Siqueira de Cerqueira, Anayran Pinheiro De Azevedo, Heloise Acco Tives Leão, and Edna Dias Canedo. 2022. Guide for Artificial Intelligence Ethical Requirements Elicitation - RE4AI Ethical Guide. In *55th Hawaii International Conference on System Sciences, HICSS 2022, Virtual Event / Maui, Hawaii, USA, January 4-7, 2022.* ScholarSpace, http://hdl.handle.net/10125/80015, 1–10.

[5] Djellel Difallah, Elena Filatova, and Panos Ipeirotis. 2018. Demographics and Dynamics of Mechanical Turk Workers. In *Proceedings of the Eleventh ACM International Conference on Web Search and Data Mining* (Marina Del Rey, CA, USA) *(WSDM '18)*. Association for Computing Machinery, New York, NY, USA, 135–143. https://doi.org/10.1145/3159652.3159661

[6] European Parliament and Council of the European Union. 2024. Regulation (EU) 2024/1689: Laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng. Published in the Official Journal of the European Union on 12 July 2024; entered into force on 1 August 2024.

[7] Isaias Ghebrehiwet, Nazar Zaki, Rafat Damseh, and Mohd Saberi Mohamad. 2024. Revolutionizing personalized medicine with generative AI: a systematic review. *Artificial Intelligence Review* 57, 5 (April 2024), 1–41. https://doi.org/10.1007/s10462-024-10768-5

[8] Clendson Domingos Gonçalves, Eduardo de Paoli Menescal, Fábio Lúcio Lopes de Mendonça, and Edna Dias Canedo. 2024. Trust in AI: Perspectives of C-Level Executives in Brazilian Organizations. In *Proceedings of the XXIII Brazilian Symposium on Software Quality, SBQS 2024, Salvador, Bahia, Brazil, November 5-8, 2024*, Ivan Machado, José Carlos Maldonado, Tayana Conte, Edna Dias Canedo, Johnny Marques, Breno Bernard Nicolau de França, Patrícia Matsubara, Davi Viana, Sérgio Soares, Gleison Santos, Larissa Rocha, Bruno Gadelha, Rodrigo Pereira dos Santos, Ana Carolina Oran, and Adolfo Gustavo Serra Seca Neto (Eds.). ACM, https://doi.org/10.1145/3701625.3701654, 147–157. https://doi.org/10.1145/3701625.3701654

[9] Tricia A. Griffin, Brian P. Green, and Jos V.M. Welie. 2024. The ethical wisdom of AI developers. *AI and Ethics* 5, 2 (March 2024), 1087–1097. https://doi.org/10.1007/s43681-024-00458-x

[10] Tricia A. Griffin, Brian Patrick Green, and Jos V. M. Welie. 2023. The ethical agency of AI developers. *AI and Ethics* 4, 2 (Jan. 2023), 179–188. https://doi.org/10.1007/s43681-022-00256-3

[11] Arif Ali Khan, Muhammad Azeem Akbar, Mahdi Fahmideh, Peng Liang, Muhammad Waseem, Aakash Ahmad, Mahmood Niazi, and Pekka Abrahamsson. 2022. AI Ethics: An Empirical Study on the Views of Practitioners and Lawmakers. arXiv:2207.01493 [cs.CY] https://arxiv.org/abs/2207.01493

[12] Fernando Korn Malerbi, Luis Filipe Nakayama, Robyn Gayle Dychiao, Lucas Zago Ribeiro, Cleva Villanueva, Leo Anthony Celi, and Caio Vinicius Regatieri. 2023. Digital Education for the Deployment of Artificial Intelligence in Health Care. *Journal of Medical Internet Research* 25 (June 2023), e43333. https://doi.org/10.2196/43333

[13] Pushkar Mishra, Helen Yannakoudakis, and Ekaterina Shutova. 2021. Modeling Users and Online Communities for Abuse Detection: A Position on Ethics and Explainability. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (Eds.). Association for Computational Linguistics, Punta Cana, Dominican Republic, 3374–3385. https://doi.org/10.18653/v1/2021.findings-emnlp.287

[14] Hamad Odhabi and Maurice Abi-Raad. 2024. Comparative Analysis of Microsoft and Google's Strategies in the Era of Advanced Artificial Intelligence Technologies. In *IBIMA Conference on Artificial intelligence and Machine Learning*. Springer, https://doi.org/10.1007/978-3-031-79086-7_4, 30–43.

[15] Aastha Pant, Rashina Hoda, Simone V. Spiegler, Chakkrit Tantithamthavorn, and Burak Turhan. 2024. Ethics in the Age of AI: An Analysis of AI Practitioners' Awareness and Challenges. *ACM Trans. Softw. Eng. Methodol.* 33, 3, Article 80 (March 2024), 35 pages. https://doi.org/10.1145/3635715

[16] Aastha Pant, Rashina Hoda, Chakkrit Tantithamthavorn, and Burak Turhan. 2024. Ethics in AI through the practitioner's view: a grounded theory literature review. *Empir. Softw. Eng.* 29, 3 (2024), 67. https://doi.org/10.1007/S10664-024-10465-5

[17] Daniel Porto, Renata Prado, Gilmar Marques, André Serrano, Fabio Mendonça, and Edna Canedo. 2025. Ethical Requirements in the Age of Artificial Intelligence: A Systematic Literature Review. In *Anais do XXI Simpósio Brasileiro de Sistemas de Informação* (Recife/PE). SBC, Porto Alegre, RS, Brasil, 663–672. https://doi.org/10.5753/sbsi.2025.246613

[18] Inioluwa Deborah Raji, Timnit Gebru, Margaret Mitchell, Joy Buolamwini, Joonseok Lee, and Emily Denton. 2020. Saving Face: Investigating the Ethical Concerns of Facial Recognition Auditing. In *AIES '20: AAAI/ACM Conference on AI, Ethics, and Society, New York, NY, USA, February 7-8, 2020*, Annette N. Markham, Julia Powles, Toby Walsh, and Anne L. Washington (Eds.). ACM, https://doi.org/10.1145/3375627.3375820, 145–151. https://doi.org/10.1145/3375627.3375820

[19] Thibault Schrepel and Jason Potts. 2025. Measuring the openness of AI foundation models: competition and policy implications. *Information & Communications Technology Law* 34, issue 3 (2025), 1–26.

[20] Senado Federal. 2023. Projeto de Lei n° 2338, de 2023: Dispõe sobre o uso da Inteligência Artificial. https://legis.senado.leg.br/sdleg-getter/documento?dm=9347622&ts=1742240889313&disposition=inline. Autoria: Senador Rodrigo Pacheco (PSD/MG); Nº na Câmara: PL 2338/2023.

[21] The White House. 2023. Executive Order 14110: Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence. https://www.federalregister.gov/documents/2023/11/01/2023-24283/safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence. Executive Order 14110, issued October 30, 2023.

[22] Rich Timpone and Yongwei Yang. 2018. Justice Rising - The Growing Ethical Importance of Big Data, Survey Data, Models and AI. In *BigSurv 18 (Big Data Meet Survey Science) conference*. Research Google, https://research.google/pubs/justice-rising-the-growing-ethical-importance-of-big-data-survey-data-models-and-ai/, 1.

[23] Ville Vakkuri, Kai-Kristian Kemell, Marianna Jantunen, and Pekka Abrahamsson. 2020. "This is Just a Prototype": How Ethics Are Ignored in Software Startup-Like Environments. In *Agile Processes in Software Engineering and Extreme Programming*, Viktoria Stray, Rashina Hoda, Maria Paasivaara, and Philippe Kruchten (Eds.). Springer International Publishing, Cham, 195–210.

[24] Ville Vakkuri, Kai-Kristian Kemell, Joel Tolvanen, Marianna Jantunen, Erika Halme, and Pekka Abrahamsson. 2022. How Do Software Companies Deal with Artificial Intelligence Ethics? A Gap Analysis. In *Proceedings of the 26th International Conference on Evaluation and Assessment in Software Engineering* (Gothenburg, Sweden) *(EASE '22)*. Association for Computing Machinery, https://doi.org/10.1145/3530019.3530030, 100–109.

[25] Ville Vakkuri, Kai-Kristian Kemell, Joel Tolvanen, Marianna Jantunen, Erika Halme, and Pekka Abrahamsson. 2022. How Do Software Companies Deal with Artificial Intelligence Ethics? A Gap Analysis. In *Proceedings of the 26th International Conference on Evaluation and Assessment in Software Engineering* (Gothenburg, Sweden) *(EASE '22)*. Association for Computing Machinery, New York, NY, USA, 100–109. https://doi.org/10.1145/3530019.3530030

[26] David Gray Widder, Derrick Zhen, Laura Dabbish, and James Herbsleb. 2023. It's about power: What ethical concerns do software engineers have, and what do they (feel they can) do about them?. In *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency* (Chicago, IL, USA) *(FAccT '23)*. Association for Computing Machinery, New York, NY, USA, 467–479. https://doi.org/10.1145/3593013.3594012

[27] Claes Wohlin, Per Runeson, Martin Höst, Magnus C. Ohlsson, Björn Regnell, and Anders Wesslén. 2012. *Experimentation in Software Engineering*. Springer, https://doi.org/10.1007/978-3-642-29044-2. https://doi.org/10.1007/978-3-642-29044-2

[28] Abdullahi Yusuf, Nasrin Pervin, Marcos Román-González, and Norah Md Noor. 2024. Generative <scp>AI</scp> in education and research: A systematic mapping review. *Review of Education* 12, 2 (Aug. 2024), 1–36. https://doi.org/10.1002/rev3.3489