Ethics in AI: how software development companies in Brazil deal with the ethical implications of AI technologies

Lucas Gabriel Teixeira da Silva[1] **, Eloize Rossi Marques Seno**[1]

[1]Federal Institute of São Paulo, São Carlos - SP, Brazil

s.teixeira@aluno.ifsp.edu.br, eloize@ifsp.edu.br

Abstract. *The use of AI technologies has become increasingly common in the software development industry. Although the ethical challenges of using these technologies have been the subject of numerous academic studies, little is known about AI software development practices. Furthermore, there are currently no effective methods and frameworks to help implement ethics at a project level. Given this context, we carried out an empirical study in order to better understand how software development companies in Brazil deal with the ethics of AI in practice. The results highlighted a lack of clarity in standard guidelines and a lack of concern about the ethical implications imposed by the use of AI technologies, suggesting that the governance of AI systems based on principled ethical guidelines is not enough to establish norms to the AI industry and its developers.*

1. Introduction

Artificial Intelligence - AI systems (or intelligent systems[1]) have become ubiquitous in our daily lives in a wide range of applications. Examples are product recommendation systems as seen in e-commerce giants like Amazon[2] and Mercado Livre[3]; fraud detection algorithms for credit cards [Ileberi et al. 2022]; stock market prediction algorithms [Bansal et al. 2022]; algorithms that analyze consumer sentiment about certain products/services on e-commerce websites [Kaur and Sharma 2023] and facial recognition security systems used to allow access to smartphones, residential condominiums or even for surveillance and public safety purposes [Silva and Silva 2019].

Despite the benefits that these AI techniques have brought to different areas of society, from marketing to healthcare, several researches have pointed out the many ethical challenges inherent to the use of AI technologies [Piteira et al. 2019, Silva and Silva 2019, Garcia 2020, Vakkuri et al. 2020, Carvalho et al. 2022]. As intelligent systems are largely software, a few ethical implications pose dilemmas for the software engineering. Some of them are technical and arise from the apparent impenetrable complexity of AI technology. Examples are the explainability (transparency) and fragility of some machine learning models as well as bias in training datasets [Larson et al. 2016, Denning and Denning 2020, Garcia 2020]. Others are social dilemmas and arise from the struggling on resolving conflicts of emotional values to the satisfaction of society as a whole, such as surveillance capitalism, employment, job, the use of AI for military purposes (e.g. AI-controlled drones capable of detecting targets

[1]Throughout this paper, we will use terms such as intelligent algorithms, intelligent systems, AI techniques and AI systems as synonyms.

[2]`https://www.amazon.com.br/` (Accessed on: May 11, 2023).

[3]`https://www.mercadolivre.com.br/` (Accessed on: May 11, 2023).

and shooting them) and for decision-making instead of humans (e.g. self-driving cars) [Denning and Denning 2020, Carvalho et al. 2022].

As algorithms take over tasks previously performed by humans, of socio-cognitive dimensions, they inherit several social responsibilities. For this reason, various studies argue that it is crucial to effectively address ethical issues of AI systems *By design*, so that ethical principles are present from their inception [Piteira et al. 2019, Vakkuri et al. 2019, Rossetti and Angeluci 2021]. Moreover, according to [Piteira et al. 2019], as important as developing powerful and scalable AI systems is to develop transparent software that can be audited and are robust enough to prevent improper manipulation.

Although there are codes of ethics with general guidelines for professional software developers, such as the documents proposed by ACM[4] and IEEE[5], and there is an extensive literature indicating the ethical aspects that must be taken into account when developing AI systems (e.g. [Larson et al. 2016, Etzioni and Etzioni 2017, McNamara et al. 2018, Piteira et al. 2019, Vakkuri et al. 2019, Garcia 2020, Rossetti and Angeluci 2021, Corrêa et al. 2022], the gap between research and practice in the area remains a challenge, as revealed by the case study carried out by [McNamara et al. 2018], that indicates the lack of importance of these ethical codes in ethical decision making.

In Brazil, specifically, little is known about software development practices in the context of AI implications, as empirical research in the area are currently scarce. In addition, due to the lack of regulation in software developer profession, studies such as [Corrêa et al. 2022] suggest that the existing development codes of ethics do not have effect on the decision-making of developers and stakeholders, since these practices are not supervised by any regulatory body.

Motivated by understanding better how software development companies in Brazil deal with ethical implications of AI technologies, we conducted an empirical research with six companies and, through a questionnaire, we collected data related to common practices used during and after the software development process. To the best of our knowledge, this is the first study in Brazil that investigated in practice how companies deal with the ethics of AI.

The rest of this paper is organized as follows. In section 2 we present the theoretical background of this research, reviewing some related work in the area. Then, in section 3, we address the elaboration of the questionnaire used for data collection in software development companies. In section 4 we detail the empirical study carried out in the companies and the results observed from the application of the questionnaire. In section 5 we present a discussion about the results and, finally, in section 6 we present the main conclusions of this work.

2. Theoretical Foundations

Ethics is a philosophical discipline that goes beyond the analysis of human behavior and focuses on the study of morality and reflection about what is good and evil in each context

[4] `https://www.acm.org/code-of-ethics` (Accessed on: May 11, 2023).

[5] `https://www.ieee.org/about/corporate/governance/p7-8.html` (Accessed on: May 11, 2023).

of society. Ethics is not static or rigid, but rather a resolution of different parameters defined by social and ethnic scopes. There has not been a single definition of ethics that has remained untouched over time, due to the evolution of morality and flexible standards [Carrillo 2020]. In this way, it is crucial to consider not only the technical aspects of intelligent systems but also their ethical implications, as they can deeply effect individuals and society. This embraces examining how an intrinsically human characteristic, such as ethical decision-making, influences the applicability of these tools.

As pointed out by [Denning and Denning 2020], algorithms can present ethical implications in many distinct contexts and different levels of severity, as exemplified by the case of COMPAS software, developed by Northpointe[6] in the United States. COMPAS is a risk assessment system used in parole hearings, which assigns scores to defendants based on questions and criminal records. However, a racial bias was identified in the software, with black defendants more likely to be misclassified as high risk compared to white defendants. Even in the face of these findings, the company maintained anti-transparency policies and continued to operate in several states [Larson et al. 2016]. These cases of racial bias cannot be justified by the removal of racial data alone, as they reflect public policies and structural inequalities present in various sectors of society [Obermeyer et al. 2019]. Ensuring the transparency of these algorithms is fundamental not only as an act of good faith by companies, but also to protect the rights of the people affected by these tools [Piteira et al. 2019]. The spread of a utilitarian ideology, which seeks the good of the majority at the expense of marginalized minorities, often perpetuates injustices under the guise that mistakes can occur naturally.

The various real-world ethical incidents caused by AI systems reported in the literature (e.g. [Larson et al. 2016, Obermeyer et al. 2019, Garcia 2020, Borges and FILÓ 2021] reinforce the relevance of being concerned with the ethical implications of these systems since their conception. An empirical study conducted by [Vakkuri et al. 2020] with three Finnish startups, aiming at understanding how the ethical implications of AI are faced in software development practices, revealed a complete lack of knowledge and concern about AI ethics. The study also showed that while some common practices, such as software documentation, can help address some ethical aspects of AI, there are no effective methods and practices to help implement ethics at the project level.

Despite the ethics of AI has been the subject of numerous academic studies in recent years, which has led to several government entities, scientific societies and companies to propose legal norms and documents with guidelines to address ethical issues during the development of AI software (e.g. Bill 21/2020[7], IEEE guidelines for Ethically Aligned Design - EAD[8], UNESCO's Recommendation on the Ethics of Artificial Intelligence[9], EU guidelines on ethics [Madiega 2019], Google's AI Responsibility Prin-

[6] `https://www.equivant.com/northpointe-suite-case-manager/` (Accessed on: Jun 29, 2023).

[7] `https://www.camara.leg.br/propostas-legislativas/2236340` (Accessed on: May 15, 2023, in Brazilian Portuguese).

[8] `https://standards.ieee.org/industry-connections/ec/ead-v1/` (Accessed on: May 16, 2023).

[9] `https://unesdoc.unesco.org/ark:/48223/pf0000381137` (Accessed on: May 16, 2023).

ciples[10]), there are currently few frameworks to support the implementation of ethical aspects in software development practice. An example of framework for approaching AI ethics through practice was proposed by [Vakkuri et al. 2019]. The framework, known as ART, is anchored by three main principles that form the basis of ethical AI development (Figure 1): accountability, responsibility and transparency, the so-called ART principles. In order to make these principles tangible, the framework proposes a set of actions (Figure 1 (1.1-3.5)) related to each one. These actions were outlined based on the IEEE guidelines EAD, which were split into two categories of ethics: (i) ethics in design, i.e. software development methods to support the implementation of ethics and (ii) ethics for design, i.e. standards that ensure the integrity of developers and users.

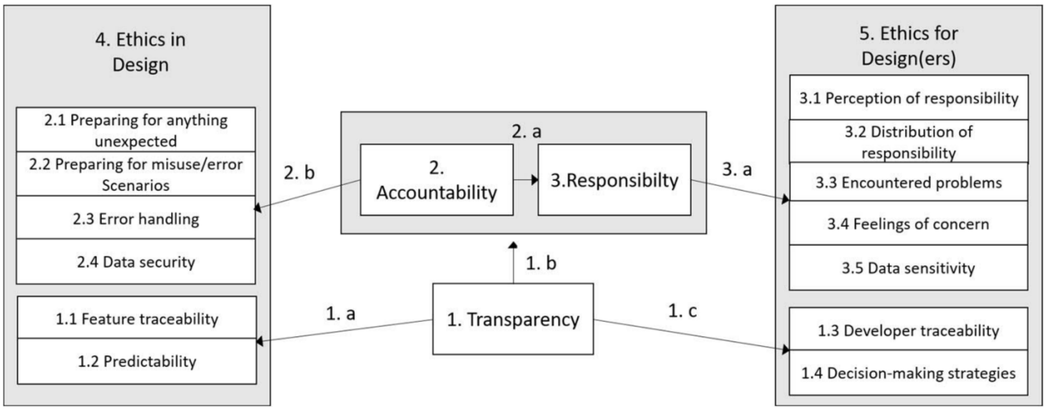

Figure 1. ART framework proposed by [Vakkuri et al. 2019]

Accountability refers to determining who is responsible for the decisions made by the software and justifying those decisions to stakeholders. The model perceives accountability through the concrete actions of developers concerning the system itself. These include preparing for unexpected situations, precaution for misuse scenarios, error handling and data security.

Responsibility is defined as an attitude or moral obligation to act ethically, which comes from an individual's internal motivation. It pertains more to the internal processes of developers rather than being directly linked to any specific behavior. Developers' actions regarding their sense of responsibility, how they distribute responsibility, handle encountered problems, and handle data sensitivity all contribute to the perception of this principle.

Transparency refers to the understanding of the internal technical processes of an intelligent algorithm. A transparent algorithm allows you to understand how the AI system generated an unexpected or erroneous output. For example, Artificial Neural Networks, the most common AI technologies, are not transparent [Denning and Denning 2020], that is, their outputs can not be easily explained.

In the ART framework, transparency is related not only to AI systems, but also

[10]https://ai.google/responsibility/principles/ (Accessed on: May 16, 2023).

to software development as a whole i.e what decisions were made, by whom, and why during the software development. According to [Vakkuri et al. 2019] some practices that supports this type of transparency are audits and software documentation. Transparency is essential for establishing accountability and responsibility. It enables the assessment of these principle by understanding how the system operates and what decisions were taken by the developers and stakeholders during development.

As the ART framework presents a practice-focused view of AI ethics, it can be used as a support tool for conducting empirical studies that aim to understand current software development practices in the context of AI, such as those carried out by [Vakkuri et al. 2020]. Inspired by [Vakkuri et al. 2020], in this work we use this framework as a basis for the elaboration of a questionnaire used to collect data in companies, as will be detailed in the following section.

3. Survey Questionnaire

Building on the ART framework and on previous research on AI ethics governance, such as [Piteira et al. 2019, Carrillo 2020, Corrêa et al. 2022], we prepared a questionnaire aiming to deepen our understanding of the principles of accountability, responsibility and transparency in companies of software development in Brazil.

We well know that, in the context of companies, the wishes and attitudes of a worker do not overlap with the orders of their superiors, often leading the worker to have unethical attitudes. With this in mind, we decided to restrict the broad concept of responsibility presented by [Vakkuri et al. 2019] and considered a more specific one focused on the professional's ethical attitudes (from now on, the principle of Professional Ethics), such as following a well-established code of ethics and making sure to do a broad and thorough risk analysis of the potential ethical implications of the product developed, taking into account the structural constraints that may limit the respondent's autonomy in these matters. Thus, we can capture the contextual limitations that shape the ethical decision-making processes within the workplace.

The questionnaire is composed of eleven questions (Q1 to Q11), including four open-ended and seven multiple-choice questions[11]. Three of them are related to principle of Transparency, four questions are about Professional Ethics, three are related to Accountability and one is a general question, which is not directly related to any of the three principles, as will be explained in Section 4.2.

4. Empirical Study

The proposed questionnaire was applied in six different software development companies. They all use AI techniques in their software processes, either as an integral part of its products or as a main product. A total of eleven responses were obtained.

The companies represented by the respondents consisted of two internationally renowned multinational corporations focused on software solutions in different areas (henceforward Case 1 and Case 2), a national holding company operating in the e-commerce sector (Case 3) and three startups, one specialized in essay automatic correction, another specialized in facial recognition solutions and another startup focused

[11]An example of the questionnaire (in portuguese) is available on `https://forms.gle/CJ8uhT7kmE6CduDQA`

on agribusiness sector (Cases 4, 5 and 6, respectively). In Case 3, we obtained multiple responses, including one from a data engineer, one from an AI analyst, and one from a project manager. For Case 6, we collected responses from one data analyst, one data scientist, and one project manager. In Case 4, we received input from two computational linguists.

The questionnaire was answered by individuals from different positions within the companies, such as one developer (9.0%), three project managers (27.2%), one data analyst (9.0%), one data engineer (9.0%), one executive director (9.0%), one AI analyst (9.0%), two computational linguists (18.1%) and one data scientist (9.0%). These respondents were categorized into three seniority levels: 45.4% were intermediate, 36.4% were junior, and 18.2% were senior. The questionnaire was designed in such a way that any level of experience could respond. Therefore, no technical questions of specific knowledge were asked, but rather questions that followed an organizational knowledge line on the part of the respondent.

The next section presents details about data collection and the section 4.2 presents the observed results.

4.1. Data collection

In order to conduct the survey, we used the Google Forms platform and sent the questionnaire directly to the companies by email. Participation in the survey was completely voluntary, meaning that employees had the option to answer the questionnaire or not. To maintain the confidentiality of the respondents and their respective companies, the responses were anonymous. We ensured the traceability of responses by sending separate copies of the questionnaire to each company that volunteered, in order to monitor the responses of each one of them.

Despite reaching out to multiple companies, only six of them kindly volunteered to respond to the questionnaire. It is important to highlight that the dissemination of the work was carried out on a voluntary basis, without offering any financial incentives for participation. We recognize that ethical topics can be sensitive to different companies, which resulted in a low response rate.

4.2. Empirical Results

The collected results are analyzed in this section according to the principles of Professional Ethics, Transparency and Accountability, as explained in Section 3. Throughout the analysis, we refer to each question in the survey (e.g., Q1, Q2), allowing the reader to better understand the observed responses. In terms of the themes explored in the questions, we can relate them to their respective actions from the ART framework depicted in Figure 1 (1.1-3.5). It is important to note that some open-ended questions can touch upon various actions or principles due to their interconnected nature, as illustrated in the framework.

Only one question (i.e. **Q4**) is not related to any of the previously mentioned principles. More specifically, it was designed to enhance our understanding of the purpose behind the utilization of AI technologies within the context of the companies under study.

We observed that the most common use of AI algorithms is related to the analysis of large sets of structured data (observed in 30.0% of the answers), in order to identify patterns and insights that can help in decision making, and large volumes of unstructured data (observed in 17.2% of cases), such as emails and social media texts. Other applications of these algorithms include creating predictive models that help forecast product trends and identify potential product failures (13.8% of cases), creating models for fraud detection (13.8% of cases), detection of anomalies in IT systems, helping to prevent failures and system interruptions (13.8% of cases) and automatizing of routine tasks (13.8% of cases). Less frequently, that is, in 3.4% of cases, AI algorithms are used for classifying and filtering products.

In the follow we present the results collected according to each analyzed principle.

4.2.1. Professional Ethics

Questions related to this principle asked respondents about the use of codes of ethics in the company focused on software development **(Q1)**, the importance of evaluating the potential impacts and risks of software during the development process **(Q2)**, the main concerns that permeate activities related to the software development process **(Q3)** and who is usually held responsible for issues related to jailbreaking or software misuse **(Q11)**. Q1 and Q2 are associated with actions 3.1-3.3 of ART framework (see Figure 1). These questions explore the definition of ethical norms, expanding the perception of responsibility, while also addressing the evaluation of risks and potential topics that need to be considered from the respondents' perspective. Q3, otherwise, is related to actions 3.4 and 3.5, as it refers to workers' concerns, such as bias, manipulation of sensitive data, among others. Finally, Q11 is linked to 3.1-3.3 due to the relation with a well-established code of ethics and it's responsibilities guidelines in case of user intervention.

Regarding **Q1**, 54.5% of respondents said that activities are guided by ethical guidelines established by the company itself, 9.0% stated that they are guided by a third-party code of ethics, with the Google's code of ethics being the most mentioned, and 36.3% stated that they are not guided by any code of ethics. Among those who said they did not follow any ethical guidelines, 75% claimed to have no knowledge of any code of ethics focused on software development, while 25% claimed to know some codes of ethics, but the company does not consider them relevant. It is important to highlight that for Case 3, that we had multiple respondents, different answers were given by people from the same team. Among the three respondents, the AI analyst indicated that the company used its own code of ethics, the project manager reported using Google's code of ethics, while the data engineer reported neither using nor being familiar with any codes of ethics that directly guide software development.

Regarding **Q2**, on a scale of 1 to 10, with 1 being not at all important and 10 being the maximum importance to analyze the possible impacts and risks of the software during the development process, 63.6% of respondents indicated the maximum importance (10), 18.2% indicated great importance (8), 9.1% indicated high importance (7) and 9.1% indicated medium importance (5). Considering all responses, the average degree of importance was 8.9.

It is worth noting that the only respondent who scored moderate importance (i.e. 5) works with facial recognition solutions, that makes this response more critical. Facial recognition tools can directly affect people if the possible impacts of these tools are not considered in advance and preventive actions are not taken, as discussed by [Larson et al. 2016], [Silva and Silva 2019] and [Garcia 2020].

With respect to the main concerns related to software development process (**Q3**, an open-ended question), the respondents' answers were quite diverse. Concerns with data security were noted in 22.7% of responses, while biases in training data sets were concerns noted in only 18.2% of answers. System reliability was also reported in 13.6% of cases and customer satisfaction in 9.1% of cases. 9.1% of cases reported they were concerned about the company not disclosing to customers that the services provided by it are performed by an AI system and not by humans. Less often, in 4.5% of cases, the respondents said concerns about system security and performance, documentation, maintenance and organization.

With respect to the accountability related to jailbreaks and misuse of software **(Q11)**, most cases, in 63.6% of answers, pointed out the project manager as the main responsible. Less frequently, respondents mentioned other responsible parties: developers (9.1%), stakeholders (9.1%), nobody (9.1%) and 9.1% were unable to answer. In Case 3, while the data engineer and project manager indicated that the team manager was responsible for jailbreak or misuse of software, the AI analyst said he is not aware of who is responsible, despite saying, in Q1, that there is an internal code of ethics in his company with well-established guidelines. In Case 6 there were also contradictory responses. While the project manager said that developers are responsible for errors arising from the software, the AI analyst and data engineer pointed out that the responsibility lies with the project manager. It is also important to say that all three interviewed mentioned (in Q1) that they are guided by ethical guidelines defined by the company itself.

4.2.2. Transparency

According to [Vakkuri et al. 2019], a system is considered transparent if the traceability feature is present and its behavior can be predicted (see Figure 1, Section 2). In addition, transparency extends beyond understanding AI systems and encompasses the entire development process, as explained in Section 2. Therefore, decisions taken in the project such as the choice of AI technology, the criteria that determine this choice, how the system will be evaluated, etc. refer to transparency.

In order to better understand the principle of transparency in the in software development practice we consider three questions about: the main criteria that determine the choice of algorithms **(Q5)**, how AI algorithms are usually evaluated **(Q6)** and how important it is to understand the chosen model in depth **(Q7)**. Q5, Q6 and Q7 are related to actions 1.1, 1.2 and 1.4. as they pertain the various aspects of transparency in the AI development process.

Regarding **(Q5)**, the following criteria were mentioned by respondents as determinants in choosing the models: performance (measured by accuracy, F1-score, etc.), in 91.0% of cases; predictability, traceability and ease of understanding of the model, in

45.4% of cases; flexibility for quick fixes and adherence to the human process of solving the same problem had an incidence of 27.3% and 18.2%, respectively. It was also evidenced here a certain degree of concern with the comprehension of algorithms, however not always followed by traceability, which could suggest the perspective of constant search for performance, encouraged by hyper-competitiveness. Because, in many cases, without a deep knowledge of the models, it is impossible to understand the reasons why the algorithm arrived at certain specific answers.

With regard to the question of how AI algorithms are usually evaluated (e.g. using classical performance measures or through user testing) **(Q6)**, 72.7% of respondents stated that they do a complete evaluation, both with classical measures and with simulations and user tests, while 27.3% of the respondents pointed out that only classical performance measure evaluations are used (e.g. accuracy, F1-score, etc.). In Case 6, the manager and the data analyst both indicated that only evaluations with classical performance measures are conducted, while the data engineer stated that they perform a comprehensive evaluation that includes both classical measures and simulations, as well as user testing.

Regarding the importance of understanding the models used in a company in depth **(Q7)**, 72.7% consider it important to understand the models in depth, while another 27.3% consider it important to understand the models just enough to be able to use it, even if they can't explain it. It is relevant to say that among those who considered it important to understand the models in depth, about 27.3% are Juniors.

Although the majority of respondents of Case 3 pointed out in **(Q7)** that they emphasized the importance of understanding the AI algorithms in depth, none of them indicated concerns about traceability in **Q5**. Instead, their responses highlighted performance, predictability of the model and flexibility of correction as key factors. This suggests that their understanding of algorithms is primarily driven by a perspective of maximizing performance, rather than an outlook for algorithm control. This observation is further supported by Case 6, where both the analyst and the data scientist pointed out that traceability is one of the criteria for algorithm evaluation. However, the project manager, in contrast to the others, emphasized the importance of ease of understanding the model and predictability as criteria.

4.2.3. Accountability

Due to the questionnaire's application context, we were able to extract a greater variety of questions related to accountability, including measures commonly adopted to prevent/treat possible data biases **(Q8)**, measures for secure data collection and handling **(Q9)** and actions taken when algorithm misbehavior is detected **(Q10)**. Q8 addresses how developers handle errors related to biases in training data, whether from a preventive or corrective perspective. This question is associated with actions 2.1-2.3 of the ART framework. Q9 is related to action 2.4 as it is focused on data security measures. Finally, Q10 is linked to actions 2.1-2.3, because of it's correlation with error handling of algorithm misbehavior in a normative design outlook. Q8, Q9 and Q10 are all open-ended questions.

With regard to bias prevention/treatment measures **(Q8)**, 25.0% of the cases reported the use of data balancing techniques and 18.8% said they made a careful data selection, aiming to ensure greater diversity and maximizing the variety of examples. Another 18.4% reported performing empirical tests using different data sets, while 12.5% claimed to use adversarial data, that is, with biases. 6.3% said they make use of whitelists and blacklists to fix known problems that an algorithm could not solve and another 6.3% they had no knowledge. Among this 6,3% that were unaware of the measures taken, one respondent from Case 3 stood out, as they emphasized the importance of a comprehensive understanding of the company's employed model (as mentioned in **Q7**). However, when asked about bias prevention and treatment measures, the respondent indicated a lack of knowledge in this regard.

When asked about the issue of data security **(Q9)**, the use of measures to restrict access to data was observed in 50.0% of responses. The treatment of sensitive data (anonymization) was mentioned in 33.0% of cases. In 8.3% of cases, we observed that there are no actions to ensure the safe collection and treatment of data and another 8.3% were unable to respond. Although issues related to data security are guided by the a Brazilian governance general data protection law[12] - LGPD, that requires the safe collection and storage of data, as well as the treatment of sensitive data and measures to restrict access to data, only one respondent directly mentioned it.

In relation to what is usually done when a bad behavior of the algorithm is detected (**Q10**), the respondents mentioned different approaches and some of them pointed out more than one: data analysis to detect anomalies (11.8%), parameter adjustment (17.6%), retraining (29.4%), training using a different model (11.8%), use of heuristics to minimize errors (5.9%), manual solutions (5.9%) and search for new features (5.9%).

5. Discussion of Results

Analyzing the answers collected from the point of view of Ethics in Design (as shown in Figure 1, Section 2), we observed that some actions that contribute to the implementation of ethics during software development do not seem to have much practical relevance in the studied companies. For example, with respect to data security (action 2.4), measures such as restricting access to data were mentioned in only 50.0% of the cases, while the treatment of sensitive data stood out in 33.0%(according to answers to Q9). This is supported by responses given to Q3 (an open-ended question), which showed that data security is not a common concern during development, been mentioned in only 22.7% of the cases.

Regarding biases (actions 2.1-2.3), although the majority of respondents claimed to use some kind of measure to avoid biases in training data (in Q8), this ethical issue was mentioned as a real concern by only 18.2% (in Q3), which according to [Garcia 2020] can turn into an issue because, algorithms tend to reproduce biases, and many times, it is not only about the omission of sensitive data, because in cases like [Obermeyer et al. 2019] and [Larson et al. 2016] contextual factors of the application can have a key role in algorithm's behavior. Finally, with regard to traceability and predictability aspects of AI algorithms (actions 1.1-1.2), these were considered important to determine the choice of AI algorithm in only 45.4% of cases, according to answers to Q5.

[12] `https://www.planalto.gov.br/ccivil\_03/\_ato2015-2018/2018/lei/l13709.htm` (Accessed on: May 18, 2023).

In a similar way, when we analyze the answers collected from the point of view of Ethics for Design(ers) (see Figure 1), some actions that ensure the integrity of developers do not seem to have much importance in practice. For example, some practices such as software documentation, that according to [Piteira et al. 2019] is crucial for enabling developer traceability and understanding the decisions made by the development team (actions 1.3 and 1.4), were mentioned as a concern only in 4.5% of cases (in Q3). In addition, although most respondents stated that they are guided by a code of ethics with well-established standards of behavior (in Q1), the results showed that in some cases there is a lack of clarity regarding the perception and distribution of responsibilities (actions 3.1 and 3.2), which can be observed in the contradictions of answers given to Q11 by the volunteers of Cases 3 and 6, corroborating for what was found by [McNamara et al. 2018] regarding the effectiveness of the recommendation ethical codes (e.g. IEEE and ACM). And despite many respondents said they consider it essential to assess the risks and potential impacts of software during development process, this concern was not mentioned by anyone in Q3, supporting the unconcerned view of the implementation of AI algorithms present in [Vakkuri et al. 2020].

With regard to feelings of concern (action 3.4), key ethical aspects such as bias, system transparency, predictability of model behavior and transparency of software development are not among the main concerns mentioned by respondents, according to the answers given to Q3. It is important to note that, open-ended questions such as Q3, Q8, Q9 and Q10, may not cover all issues that developers and stakeholders consider important in practice.

Concerning the low response rate, it is crucial to emphasize that, as discussed in Section 4.1, the limited participation in the questionnaire presents challenges in comprehensively grasping the ethical implications of AI practices in companies in Brazil. However, the employed methodology in this survey contributes to the field by offering a contextual framework for future data collection endeavors. By implementing this questionnaire in subsequent studies, it will be feasible to gather supplementary data on the topic, enabling the establishment of more robust analysis parameters on a larger and national scale. This iterative approach will enhance our understanding of the ethical landscape surrounding AI and foster more informed decision-making in this domain.

Comparing the results obtained by [Vakkuri et al. 2020] from Finnish startups, Brazil reveals an organizational challenge. Beyond the lack of concern regarding the potential impacts of intelligent algorithms, respondents highlighted structural issues in security and decision-making. These arose due to the absence of traceability and guidelines for some software maintenance scenarios. This reinforces the notion that development practices should adopt a normative perspective to influence software development. This becomes especially pertinent given the ineffectiveness of ethical recommendations [McNamara et al. 2018]. Without this normative approach, these development practices can potentially lead to issues in critical sectors (i.e. [Obermeyer et al. 2019] and [Larson et al. 2016]). These situations can consequently introduce further ethical dilemmas associated with AI accountability as seen in [Denning and Denning 2020].

6. Conclusion

This paper presents an empirical study that aimed at understanding how software development companies that use AI technologies have dealt in practice with the ethical implications of AI. Based on the ART framework proposed by [Vakkuri et al. 2019] and on previous research on AI ethics governance (e.g. [Piteira et al. 2019, Carrillo 2020, Rossetti and Angeluci 2021, Corrêa et al. 2022]), we proposed a questionnaire to be applied to AI software development companies in Brazil. We examined the current state of ethics implementation, considering the lack of laws and responsible regulatory bodies in Brazil.

The observed results showed that codes of ethics and professional conduct such as those proposed by scientific societies like ACM and IEEE have limited influence on the decision-making of developers and stakeholders due to most of the companies studied reported the adoption of ethical guidelines established by themselves. As a consequence, it is highlighted that there is little concern with the ethical implications of using AI technologies such as biases, the risks and potential impacts of the software, lack of transparency both in the functioning of AI systems and in the decision-making process during software development. These findings corroborates the study done by [McNamara et al. 2018] and [Corrêa et al. 2022], which suggests that this form of governance based on principled ethical guidelines might not be enough to normalize the AI industry and its developers.

Further studies on this topic should aim to expand the application of the questionnaire within the Brazilian context and focus on developing regulatory research that bridges the gap between ethical recommendations and ethical development practices.

References

Bansal, M., Goyal, A., and Choudhary, A. (2022). Stock market prediction with high accuracy using machine learning techniques. *Procedia Computer Science*, 215:247–265.

Borges, G. and FILÓ, M. (2021). Inteligência Artificial, Gênero e Direitos Humanos: O Caso Amazon. *Revista Justiça do Direito*, 35(3):218–243.

Carrillo, M. R. (2020). Artificial Intelligence: From ethics to law. *Telecommunications Policy*, 44(6):101937.

Carvalho, L. P., Murakami, L., Suzano, J. A., Oliveira, J., Revoredo, K., and Santoro, F. M. (2022). Ethics: What is the research scenario in the brazilian conference bracis? In *Anais do XIX Encontro Nacional de Inteligência Artificial e Computacional*, pages 624–635. SBC.

Corrêa, N. K., de Oliveira, N. F., and Massmann, D. F. (2022). Sobre a Eficiência da ética como Ferramenta de Governança da Inteligência Artificial. *Veritas (Porto Alegre)*, 67(1):e42584–e42584.

Denning, P. J. and Denning, D. E. (2020). Dilemmas of Artificial Intelligence. *Communications of the ACM*, 63(3):22–24.

Etzioni, A. and Etzioni, O. (2017). Incorporating Ethics into Artificial Intelligence. *The Journal of Ethics*, 21:403–418.

Garcia, A. C. (2020). Ética e inteligencia artificial. *Computação Brasil*, pages 14–22.

Ileberi, E., Sun, Y., and Wang, Z. (2022). A machine learning based credit card fraud detection using the ga algorithm for feature selection. *Journal of Big Data*, 9(1):1–17.

Kaur, G. and Sharma, A. (2023). A deep learning-based model using hybrid feature extraction approach for consumer sentiment analysis. *Journal of Big Data*, 10(1):1–23.

Larson, J., Mattu, S., Kirchner, L., and Angwin, J. (2016). How We Analyzed The Compas Recidivism Algorithm, 2016. *URL https://www. propublica. org/article/how-we-analyzed-the-compas-recidivism-algorithm*.

Madiega, T. A. (2019). Eu guidelines on ethics in artificial intelligence: Context and implementation.

McNamara, A., Smith, J., and Murphy-Hill, E. (2018). Does ACM's Code of Ethics Change Ethical Decision Making in Software Development? In *Proceedings of the 2018 26th ACM joint meeting on european software engineering conference and symposium on the foundations of software engineering*, pages 729–733.

Obermeyer, Z., Powers, B., Vogeli, C., and Mullainathan, S. (2019). Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations. *Science*, 366(6464):447–453.

Piteira, M., Aparicio, M., and Costa, C. J. (2019). Ethics of Artificial Intelligence: Challenges. In *14th Iberian Conference on Information Systems and Technologies (CISTI)*, pages 1–6. IEEE.

Rossetti, R. and Angeluci, A. (2021). Ética algorítmica: questões e desafios éticos do avanço tecnológico da sociedade da informação. *Galáxia (São Paulo)*.

Silva, R. L. d. and Silva, F. d. S. R. d. (2019). Reconhecimento facial e segurança pública: os perigos do uso da tecnologia no sistema penal seletivo brasileiro. In *Congresso Internacional de Direito e Contemporaneidade, Santa Maria, RS, Brasil*, volume 5.

Vakkuri, V., Kemell, K.-K., and Abrahamsson, P. (2019). AI Ethics in Industry: A research framework. In *Proceedings of the Third Seminar on Technology Ethics - Tethics*, pages 49–60.

Vakkuri, V., Kemell, K.-K., Jantunen, M., and Abrahamsson, P. (2020). This is Just a Prototype": How ethics are ignored in software startup-like environments. In *Agile Processes in Software Engineering and Extreme Programming: 21st International Conference on Agile Software Development, XP 2020, Copenhagen, Denmark, June 8–12, 2020, Proceedings 21*, pages 195–210. Springer International Publishing.