World Journal of Advanced Research and Reviews

eISSN: 2581-9615 CODEN (USA): WJARAI

Cross Ref DOI: 10.30574/wjarr

Journal homepage: https://wjarr.com/

(Review Article)

Existing challenges in ethical AI

Sameerkumar Babubhai Prajapati *

Computer Science, Judson University, USA.

World Journal of Advanced Research and Reviews, 2025, 25(01), 2130-2137

Publication history: Received on 16 December 2024; revised on 23 January 2025; accepted on 26 January 2025

Article DOI: https://doi.org/10.30574/wjarr.2025.25.1.0267

Abstract

The current growth of Artificial Intelligence (AI) is enormous, and has the potential of changing industries, economies and societies in the today's world. Nevertheless, as AI is more and more involved in decision-making processes, the ethical issues come out and affect trust, fairness, and accountability of these systems. From the reviewed literature, this paper identifies five key ethical issues in AI: Algorithmic Bias, Fairness, Transparency, Accountability, and Privacy. Nevertheless, the current AI systems have a tendency to mimic the prejudices which have existed throughout history, which presents threats to society equity. Furthermore, a question of accountability and explanation of every decision made by AI brings the question of the explainability of AI and the ability to track the reasoning of an AI model. There are also issues to do with privacy as much of the data needed for training AI models is personal data and this is a major issue of privacy. The paper elaborates these challenges and recommends the following solutions: design non-discriminatory and accountable AI systems; follow ethical principles in AI development; test AI systems for biases; and encourage interdisciplinary cooperation. The paper also seeks to discuss the most effective ways of incorporating ethical principles into the AI development process and supports the element of regulation to enhance the existing rules. Through addressing these significant concerns, this paper seeks to contribute theoretical and practical recommendations towards the proper creation and application of AI-based solutions for the collective benefit of society and creating confidence in the system.

Keywords: Artificial Intelligence (AI); Ethical AI; Algorithmic bias; Fairness; Bias mitigation; Fairness constraints; Adversarial debiasing; Data augmentation; Transparency; Accountability; Privacy; Machine learning; Hiring systems; Demographic equity; Ethical frameworks

1. Introduction

1.1. Context and Importance of AI

Elements of artificial intelligence, from a theoretical concept, have become a crucial part of the modern society. The roots of it began in early research in the mid-20th century, but only in recent decades has it made significant strides in a variety of applications. Currently, in healthcare, AI technologies are being used to improve diagnostics, discover drugs, and improve personalized treatment plans to ultimately help reach a better patient outcome. In financial terms, AI is used to detect the risk, the fraud and even to trade, indirectly, faster, more efficiently, and cheaper. AI tools are being used in law enforcement to predict crime patterns, carry out surveillance or even conduct investigations. These applications represent just some of the huge potential of AI to reduce bottlenecks in the system, increase productivity and solve complex problems in areas that humans cannot solve by themselves.

* Corresponding author: Sameerkumar Prajapati

1.2. AI in Modern Society

Although usually associated with industry specific applications, AI are creeping their way into everyday life. Whether it's voice assistants like Siri and Alexa, or AI driven recommendation systems for Netflix and Amazon, we live in a world that's riddled with AI. AI tools are augmenting tasks, automating decisions or generally having a positive impact in the workplace. However, with this massive adoption of AI, there are critical unknowns about how AI will alter work—what jobs can be automated, which skills will be required, and how ethical it can be to automate.

More importantly, AI is involved more actively in decisions that directly impact people's lives, including loan approvals, criminal justice sentencing, and who gets a job. Increasingly, AI is playing a key role in critical decisions, and fairness, bias and accountability are being called into question. AI systems are becoming more autonomous, more sophisticated, and therefore taking on a greater number of higher stakes: It serves to point the direction of entire industries, of economies, of societies.

1.3. Need for Ethical AI

Therefore, its embedding of ethical considerations in AI technologies is essential. Fair AI, transparent AI, inclusive AI: These are only a few of a number of terms that are used to ensure that these systems work in a manner that is fair, open, doesn't favor one person over another and that everyone gets a shot. Without such frameworks, AI systems can entrench bias, trample privacy and exercise opaque decision making that mutates public trust. This means that developers, researchers, and policymakers alike have this ethical issue to solve, effectively on the front foot, ensuring that potential AI benefits us all without perpetuating inequalities or creating new divides in society.

1.4. Problem Statement

With AI impacting more and more parts of our lives, we can also see the ethical challenges that come with that rapid growth at hand. Key issues with regard to bias in algorithms, lack of transparency and accountability in the decision-making processes are prevalent among both those using AI and those developing AI. The danger of these ethical challenges is, quite simply, that they threaten to undermine the positive potential of AI and so need addressing head on. We explore the ethical challenges of existing efforts to develop and deploy ethical AI in this paper, identify the most pressing issues and outline potential solutions to promote responsible design and use of AI.

Table 1 Key Areas Affected by AI in Different Sectors

Sector	Key Applications	Ethical Concerns
Healthcare	Diagnostic tools, robotics	Bias, privacy
Finance	Credit scoring, trading	Fairness, accountability
Criminal Justice	Predictive policing	Bias, transparency

2. Key Ethical Challenges in AI

2.1. Algorithmic Bias

2.1.1. Definition

Bias in AI is a problem when data or even algorithms give preferences to one set of people over another. This bias can be explicit or implicit, and usually correspond to something being loaded in the data that we use to train AI models. Because AI systems are frequently built to maximize based on trends in data, placenta any inequity in that data can yield inequity in their output, causing the inequity to persist.

2.1.2. Manifestation of Bias

From hiring algorithms to predictive policing tools, AI bias shows up in many different industries. When our AI models are trained on historical hiring data, that data can be polluted with past bias in the workforce. Predictive policing tools that rely on historical crime data may do the same, unnecessarily further the harm of biased policing against minority communities in particular. Also, although credit scoring models are based on personal data, they can also inadvertently amplify discrimination and exacerbate the economic exclusion and opportunities of some.

2.1.3. Real-World Examples

- Hiring Algorithms: This means that many of the available AI driven recruitment tools favor male candidates or certain ethnic groups based on the datasets that are used to train them (which historically favored such candidates in hiring decisions) (Chen, Storchan, & Kurshan, 2021). It's an infamous example — Amazon scrapped an AI recruitment tool for being biased against women. Trained on resumes from Amazon candidates, which overwhelmingly hail from the tech industry that is mostly male, the tool learned to show preference to male candidates.
- Predictive Policing: Predictive policing algorithms have been accused of boosting racial profiling. For example, risk assessment tools could be used in law enforcement as a means of over police minority communities. For example, in the U.S. COMPAS (Correctional Offender Management Profiling for Alternative Sanctions), which has been demonstrated to be prone to overprediction of Black defendants' risk as high as compared to white defendants (Khanna et al., 2020).
- Credit Scoring Models: AI driven credit scoring models in financial services can exacerbate preexisting economic inequalities. Historical financial data is often applied to these models, which can undermine historically marginalized groups without access to credit, including low income and minority communities (Kurshan et al., 2021).

2.1.4. Challenges

The main problem with algorithmic bias is the quality of the data with which we train AI models. In other words, if the data used to train the system is unrepresentative or skewed in a specific way, the AI will pick them up and grow up to adapt them into their behavior. However, the sector lacks the tools and frameworks to detect and mitigate bias in the ways AI developers need them to. Additionally, these models are difficult to be transparent, many of these AI systems are 'black boxes', hard to understand how such biased decisions are being made.

Table 2 Examples of Algorithmic Bias in Different AI Applications

Application	Example of Bias	Impact
Hiring Algorithms	Gender and race bias in hiring	Discriminatory hiring practices
Predictive Policing	Over-policing in minority areas	Heightened racial profiling
Credit Scoring Models	Discriminatory credit scoring	Economic inequality

2.2. Fairness

2.2.1. Fairness Metrics

There are multiple ways of defining, or measuring, fairness in AI. Statistics used as fairness metrics include statistical parity (same outcomes for each group), and of equalized odds (same outcome differences for each group). Unfortunately, these metrics almost always disagree with one another, and it is not possible to obtain fairness without sacrificing other goals such as accuracy or efficiency.

2.2.2. Types of Fairness

- Group Fairness: This method guarantees to treat different groups (i.e., different race or gender groups) equally or alike (in terms of treatment or outcome) by an AI system. For instance, for hiring algorithm, for group fairness, the algorithm should hire equal number of men and women for a given position.
- Individual Fairness: It means that like cases should be treated alike. For example, two applicants would be equally known for their qualification if they apply for the same job, if they have the same gender or not. Even though individual fairness might make less 'just' decisions on a day-to-day basis, it can be very hard to achieve in practice in datasets which have large amount of data.

2.2.3. Challenges

One difficulty in designing AI systems is making them fair, as well as other important goals. Fairness constraints can sometimes degrade accuracy of AI. One example is an AI model trained to promote gender balance in hiring, in which the model may turn a blind eye to the more qualified candidates in favor of gender parity. Moreover, what constitutes fairness is ambiguous; it is based on different cultures and communities who have different definitions as to what that means. Fairness is still hard to universally define.

2.3. Transparency & Accountability

2.3.1. Need for Transparency

In the world of AI, the reason why ethical issues regarding it become problematic is because decisions made by AI systems are largely opaque, not transparent. There are many AI models nowadays, especially those that based on deep learning, where the operations inside are opaqueness for everyone, from users to developers. The lack of transparency in the function also makes it difficult to verify what, and why, decisions are being made, something that might be highly imitate in sensitive applications, such as in criminal justice, healthcare and finance.

Not only do you want transparency so you know your AI system is fair and unbiased, but you also want visibility into what the AI system did in order for the user to trust it. Users that cannot understand how the AI derives at those conclusions are likely to be less willing to trust with the system, particularly the system which could have a large impact on their lives.

2.3.2. Real-World Example

When it comes to credit scoring or loan approvals in the field of financial services, AI led algorithms do not often take the time to provide an explanation of the way in which a decision was made. If the lenders use algorithms to make decisions on whether or not to lend, and they appear to be unable to challenge or even fully understand these decisions, then they could deny a consumer a loan. The lack of transparency is a big block in making sure AI systems are fair and held accountable.

Visibility in AI models is difficult, as tasks become increasingly complex while machine Learning algorithms get harder to break down. In addition, models may be made more interpretable, but even then, for example, how do we hold an AI system accountable when it makes a harmful or biasing decision? Unlike humans, AI systems make no choices; instead, the choice is made by the developers, the organizations who train these models, and the policymakers who deploy them. There is a need to create clear accountability mechanisms so as to avoid creating harm and to encourage responsible use of AI.

Table 3 Transparency and Accountability Challenges

Challenge	Description	Potential Solution
Lack of explainability	AI models can be opaque in decision-making	Implement explainable AI (XAI)
Responsibility gaps	Unclear accountability when AI makes harmful decisions	Clear regulatory frameworks

3. Possible ways forward for Ethical AI

3.1. Designing for Fairness

3.1.1. Incorporating Bias Mitigation

Incorporating your bias mitigation techniques into the development process help you address algorithmic bias. Basically, we use methods like fairness constraint, adversarial debiased, and data augmentation to eliminate bias in the training dataset and the AI model. Machine learning algorithms can be strategically designed to incorporate fairness constraints, such that predictions do not disproportionately affect specific groups. For example, constraints may be added to models to ensure that different demographic groups are treated equally in order to address disparities in outcome. Adversarial debiasing techniques, too, are similar in which we train a model so that it becomes less sensitive to biased data and directs its attentions more toward meaningful variables that do not reflect societal biases.

Alternatively, data augmentation involves the generation of additional distinct data and its inclusion in the training set to offset the biases in the original data set. It will allow us to balance the data so that the model does not learn their patterns wrong. Such techniques can be integrated into AI systems with the aim of designing fairness into them and reduce the probability of discriminatory outcomes.

3.1.2. Case Study: Algorithms for Enhancing Fairness in Hiring or in Loan Approval Systems

Take hiring or loan approval processes as an example: AI systems used here. Historically these systems have been criticized for that perpetuation of discrimination on race, gender and socioeconomic status. In a hiring system, a fairness enhancing algorithm may be one that anonymizes demographic information (for example, name or gender), to ensure that candidates apply, assess them on the basis of their qualifications, skills, and experience alone. There is a similar approach towards loan approval systems that would correct historical biases in its credit scoring models by taking in alternative data points or paying more attention to underrepresented classes to make sure the system does not deprive those who need access to loans more.

According to the AI Now Institute in 2018, when algorithms in hiring take courses in an assumption of sorts, making sense of the biased patterns in historical data, fairness metrics and techniques can be used to increase the diversity of a hiring process and provide equality. Fairness is important to creating trust, therefore, AI systems driven by fairness intended to do so can create more inclusive opportunities that are key to creating trust and promoting equity.

3.2. Explainability and Transparency

3.2.1. Explainable AI (XAI)

When working with complex machine learning models, many people consider them black boxes and they have limited transparency, which is one of the big things in AI development. Explainable AI (XAI) strives to make AI more explainable to users by providing clear, human understandable explanations of how decision-making is being impacted by the algorithm. To ensure that AI decisions are accountable, fair and that people trust AI, it is extremely important in high stakes sectors as in healthcare, criminal justice and finance.

For example, if we're talking about healthcare, an AI system will provide a treatment plan, but absent that transparency, patients and medical professionals may be hesitant to put all their eggs in its basket. XAI methods such as LIME (Local Interpretable Model Aware Explanations) and SHAP (SHapley Additive explanations) asks of model how it came to a specific output and highlights the most important features that brought about this outcome. They offer local explanations, so they explain a specific prediction made on a particular instance, not the model itself.

3.2.2. Tools and Techniques

- LIME (Local Interpretable Model-agnostic Explanations): We have been using LIME as a popular technique for producing explanations for complex models, especially in cases where interpretability is crucial. LIME approximates (locally) a model's decision process and thus it is easier to interpret. LIME, for example, could show in a loan approval model that applicant's credit score and income level were the most important factors in the decision but not much importance tied to other factors such as age or gender.
- SHAP (SHapley Additive exPlanations): Using cooperative game theory, SHAP values are a means to measure the contribution of each feature to an output of a model. Shapely values make sure the process is understandable by calculating how much each feature contributes to the prediction of a model. This means that in hiring AI systems, using SHAP it would be possible to explain why the candidate chosen for a position has been accepted, clearing any ambiguity in the decision-making process.

But these tools will significantly enhance the transparency of AI systems and allow users and regulators to grasp easily and trust AI decisions.

3.3. Enhancing Accountability

3.3.1. Accountability Mechanisms

With the role of AI increasingly intertwined with the decisions we make; we need increasingly clear accountability frameworks to both promote responsible AI development and deployment. This essentially means how we build processes for them to take responsibility for the decisions that they are making. To mitigate the potential negative impact of AI, one proposed solution is to establish AI governance boards consisting of multidisciplinary experts to oversee its ethical considerations in the development and deployment of AI systems.

Furthermore, developers are able to build in audit trails into their AI systems that log decisions, the reasons for those decisions and any changes made to the model over time. Through this, stakeholders get below decisions to its source, making them more accountable and making the detection and solving of biases, and other ethical issues, easier.

3.3.2. Regulatory Approaches

Ensuring that AI systems meet ethical standards is a huge part of regulation. But governments around the globe and even international bodies alike are starting to regulate AI technologies to defend public interests. For instance, the

European Union's General Data Protection Regulation (GDPR) includes AI transparency and accountability provisions, relevant to personal data. The requirements herein for AI developers are clear: systems should be explainable and fair.

According to Kurshan et al. (2021), global regulations on AI ethics are starting to appear. It is concluded that policy makers should acknowledge specific sectoral rules alongside the general ethical ethics of AI systems. Governments can guide the use of AI through regulatory measures like mandatory audits, certifications, and transparency standards to promote AI accountability, so that AI systems are guided by ethical principles, and are not used in exploitative or harmful ways.

3.4. Privacy Protection

3.4.1. Privacy-Preserving Algorithms

Privacy is another important ethical challenge from the AI world. Consequently, AI systems, in particular, those that are based on a big data, have the potential to amount to the infringement of individual privacy through the processing of the sensitive personal data. To tackle these problems, privacy preserving algorithms have been developed which allow AI systems to work without revealing the individual data exposed to unnecessary risk.

Federated learning is a machine learning technique that promises to solve this by allowing the models to be trained on decentralized data without sharing the raw data itself. It minimizes the likelihood of mistakes made in capturing data, and protects sensitive data from being read. Indeed, federated learning can be applied to train healthcare AI models across hospitals or institutions without transmitting patient data in order to achieve privacy as well as keep the model integrity.

We may use differential privacy, where the dataset is noisy before we use it for training. This technique makes sure that the data cannot be tracked down to any person thereby helping privacy of the users. Already, companies such as Apple and Google have rolled out differential privacy in some of their products, including iOS and Android apps to protect user data while still using it for helpful applications.

With the help of federated learning and differential privacy, developers implement privacy preserving technologies to limit the leak of privacy information of the user while they are using an AI system. That way the AI system will comply with privacy standards and respect the user's rights. As AI adoption continues, these techniques will become increasingly important to preserve this balance — we want to make use of data in service of AI, but we also want to ensure that individual privacy is preserved.

4. Case Studies

4.1. AI in Healthcare

4.1.1. Ethical Dilemmas

However there hasn't been a lot development of the AI in healthcare particularly for cancer (oncology) where it's being utilized for research, treatment planning and the prediction of the patient outcome. However, the same other side has also raised other ethical concerns about integrating AI in healthcare. One problem is privacy of the patient. With healthcare AI systems often acting on massive amounts of patient data to make predictions or recommendations they tend to raise data security and access concerns. Its crucial for AI systems to adhere to privacy regulations—such as Health Insurance Portability and Accountability Act (HIPAA) in the USA, or the European equivalent, General Data Protection Regulation (GDPR)—that protect a sick patient's personal information.

Just another ethical challenge AI models transfer in the form of treatment bias. Currently, AI systems can unintentionally keep perpetuating deeply rooted biases in healthcare and keep perpetuating unequal treatment. If algorithms have been trained on datasets that are biased (e.g. lack of minority populations), they may then diagnose less accurately or (or be) less effective for treatment for these groups. Khanna et al. (2020) explored how oncology management AI models can do much worse when trained using skewed data sets vs serving underrepresented populations in clinical trials. And it may lead to making decisions that most greatly affect some groups of people, which could increase health disparities.

4.1.2. Solutions in Practice

All these ways can be addressed. The first solution is to ensure that from the design of the AI systems to the mitigation techniques adopted aim at reducing the health disparities. We've also got some examples like training AI models with more diverse and representative datasets to ensure that appropriate demographic groups are represented in a fair manner. In addition, privacy preserving AI technologies like federated learning, which can do the processing of data decentralized, can tackle these privacy risks and still learn the AI models from large population datasets.

Along with that deployment of AI in healthcare requires monitoring and human in the loop (HITL) solutions. AI can be a support tool instead of being a replacement in the ability to include the human expertise in the decision-making process and that when predictions grow to be inaccurate or unfair, we can intervene.

4.2. AI in Finance

4.2.1. Ethical Issues

Financial sector is one of the shining examples of wide spread adoption of AI. More specifically, AI based systems for credit score and loan approval processes are becoming a place of recourse for making high stakes financial decisions. There are a lot of real ethical quandaries these systems present—most notably around fairness and transparency—but they also provide real benefits. The majority of AI credit scoring algorithms to ascertain if someone qualifies for a loan or credit rely on tons of data, from consumer spending habits to social media activity and previous history of loans, amongst other things. Efficiency can be increased, and human error decreased, but these systems can also perpetuate discrimination, unless carefully watched.

One of the more urgent of the ethical questions concerns the prospect of unfair treatment of some sections of people in lending. Take, for example, an AI model that denies people credit applications with no thought to realize an actual data point, like zip code or past employment history, is actually correlated with race or socio-economic status. It also can be disproportionately experienced by minority or low-income communities – resulting in financial exclusion and widening the wealth divide. Second, AI systems are nontransparent with respect to decision reasons for why they're denying credit to a consumer or how they made a decision given these decisions make it difficult for a consumer to fully understand the basis of the credit decision.

4.2.2. Proposed Solutions

A solution to such ethical challenges is transparency. This means that AI models that are used in financial services need to be explainable – that people have to understand why they are being made. To remove the black box nature of financial models and provide clear reasons for a credit decision, we apply explainable AI (XAI) tools such as LIME and SHAP. It also forces financial institutions to be accountable to consumers and trust AI systems.

Furthermore, financial AI models need to be made fairer both in the sense that they do not drive inequality. This means that researchers like Kurshan et al. (2021) propose that fairness in financial AI can be reached by clarifying what it means and by embedding this definition as a constraint during the development of the model and by checking regularly that a system, which implements the model, complies with ethical standards. To reduce the potential for unfair results if financial AI systems are not incorporated with diversified and representative data, as well as to make the decision-making process transparent.

AI driven systems need to be included in regulatory frameworks that include the Fair Lending Act and also need to be fair to meet ethical and legal standards. In fact, regulatory bodies might even ask for periodic versions of the audits in the financial sector that traditional regulators use – audits that would ensure that the created models do not just unintentionally reproduce discriminatory practices.

5. Conclusion

5.1. Challenges and Solutions Summary

Unsurprisingly, the integration of AI into sectors like healthcare, finance, and law enforcement comes with numerous ethical challenges. This perpetuation of historical inequality can be from algorithmic bias, which comes from unrepresented datasets, especially on things like hiring and policing. The fundamental question of fairness in AI still raises a lot of controversy because balancing the goals of fairness and other performance metrics, such as accuracy, turns out to be an unattainable balance. This is further aggravated by the 'black box' nature of many AI models, where it is hard to grasp the rationale of decisions. Finally, one is that given that AI systems work best with significant amounts of personal data, privacy concerns are exacerbated.

To mitigate these challenges, the paper proposes several solutions: Incorporating bias mitigation techniques into the design of AI systems that are fair by definition, developing means of transparency for AI (XAI), forming accountability mechanisms for developers and operators, and securing privacy with novel technologies such as federated learning and differential privacy. Furthermore, these are technical solutions that also need to be accompanied by appropriate policy and regulatory frameworks, so that they are both followed and complied with, according to ethical standards.

5.2. Implications for the Future

In the future, AI will require this ongoing collaboration across disciplines: technology, law, ethics, policy, but also data, society, and policy. If the AI developers need to be innovative, they need to work very closely with ethicists, legal experts and policymakers, in order to create the systems that are innovative but also responsible. We can only achieve this through disciplinary collaboration, first to design and then deploy AI technologies for fairness, transparency, and accountability.

5.3. Call to Action

Building an ethical AI is everyone's responsibility. It is mandatory upon the stakeholders, being AI researchers, developers and policymakers, to come up with solutions proactively to these challenges such that AI serves the broader interests of society. As AI develops, we need to think about societal impact of AI and act to develop technologically advanced and ethical systems.

References

[1] Stahl, B. C., & Stahl, B. C. (2021). Ethical issues of AI. Artificial Intelligence for a better future: An ecosystem perspective on the ethics of AI and emerging digital technologies, 35-53.

[2] Chen, J., Storchan, V., & Kurshan, E. (2021). Beyond fairness metrics: Roadblocks and challenges for ethical ai in practice. arXiv preprint arXiv:2108.06217.

[3] Kurshan, E., Chen, J., Storchan, V., & Shen, H. (2021, November). On the current and emerging challenges of developing fair and ethical AI solutions in financial services. In Proceedings of the second ACM international conference on AI in finance (pp. 1-8).

[4] Khanna, S., Srivastava, S., Khanna, I., & Pandey, V. (2020). Ethical Challenges Arising from the Integration of Artificial Intelligence (AI) in Oncological Management. International Journal of Responsible Artificial Intelligence, 10(8), 34-44.

[5] Eitel-Porter, R. (2021). Beyond the promise: implementing ethical AI. AI and Ethics, 1(1), 73-80.

[6] Belle, V. (2023). Knowledge representation and acquisition for ethical AI: challenges and opportunities. Ethics and Information Technology, 25(1), 22.

[7] Mika, N., Nadezhda, G., Jaana, L., & Raija, K. (2019). Ethical AI for the Governance of the Society: Challenges and Opportunities. In CEUR Workshop Proceedings (Vol. 2505, pp. 20-26).

[8] McKay, F., Williams, B. J., Prestwich, G., Bansal, D., Hallowell, N., & Treanor, D. (2022). The ethical challenges of artificial intelligence-driven digital pathology. The journal of pathology: clinical research, 8(3), 209-216.

[9] Rezaei, M., Pironti, M., & Quaglia, R. (2024). AI in knowledge sharing, which ethical challenges are raised in decision-making processes for organisations?. Management Decision.

[10] Al-kfairy, M., Mustafa, D., Kshetri, N., Insiew, M., & Alfandi, O. (2024, August). Ethical challenges and solutions of generative AI: an interdisciplinary perspective. In Informatics (Vol. 11, No. 3, p. 58). MDPI.