

Ethics Empowered: Ethics in Artificial Intelligence (AI)



Artificial Intelligence (AI) is transforming how we work, live and connect with each other. From automating tasks and facilitating human interactions to enhancing healthcare diagnoses and powering self-driving cars, AI offers remarkable opportunities. However, as AI becomes more integrated into industries and daily life, it is important to acknowledge and address the ethical considerations that come with it.

AI-enabled technologies are largely powered by Large Language Models (LLM) that are trained on vast amounts of data such as text, social media posts, news articles and scientific research. Such models are designed to understand, interpret and generate responses as well as perform wide range of tasks. These models can unintentionally absorb inappropriate or biased content, which may influence their outputs, thereby potentially generating harmful, biased or misleading results.

The potential for bias, privacy breaches and unforeseen consequences necessitates a proactive approach to ensure that there is responsible development and deployment of AI. To mitigate these risks, AI users need to be well-informed about the ethical principles, regulatory frameworks and best practices that guide the development, deployment and use of AI-enabled systems.

This article will explore some ethical considerations in the use of AI, bias and discrimination in AI systems, regulatory frameworks, the future of ethical AI and address ethical challenges posed by AI.

Key Ethical Considerations

As the use of AI is increasingly adopted in various industries such as manufacturing, health care, logistics and supply chain, etc., some ethical consideration such as fairness, accountability, sustainability and transparency should guide its development, deployment and use. These principles also align with legal standards like the Data Protection Act, 2012 (Act 843) and the General Data Protection Regulation (GDPR). Ethical concerns presented by AI varies according to information the AI technology receives and the responses expected. For example, a machine learning algorithm trained to detect spam emails may pose fewer ethical challenges than one trained to diagnose diseases or perform facial recognition. Each AI system should be evaluated based on its context and potential impact.

Fairness

AI systems reflect the biases of the humans who create them. Fairness in the development and deployment of AI applications means ensuring that AI tools are free of unfair bias, designed to be inclusive for diverse stakeholders and do not reinforce existing inequalities against marginalized groups. In a diverse society like Ghana, AI solutions in this environment should promote inclusivity and equity. AI systems should be trained on relevant, accurate and generalized datasets excluding variables that are ethically unreasonable. For example, a facial recognition system that is only trained on lighter-skinned individuals can lead to discriminatory outcomes. Ensuring fairness means eliminating unreasonable variables and training users to deploy AI responsibly.

Accountability

AI systems are used by various professionals to augment their experience when making decisions, whether in finance, healthcare, governance or business advisory. They can produce biased or incorrect answers, which can seriously affect the quality of decisions made. For example, if an AI-driven diagnostic tool used by a medical practitioner provides inaccurate assessments, it

could lead to misdiagnoses or wrong medical advice, potentially putting a patient's life at risk. Accountability means that users—not just the technology—are responsible for outcomes. Accountability ensures that AI remains a tool to support, not replace, human judgment. Organisations developing or using AI systems, whose outcomes may impact on people, need to carry out risk and impact assessments and put in place appropriate AI governance.

Sustainability

Developers and users of AI need to consider the long-term impacts of AI technologies on individuals, society and the environment. They must be conscious of the broader effects AI innovations may have over time. This includes considering the social, economic and environmental consequences of deploying AI systems at scale. For instance, the increasing use of AI requires significant computational power, which can contribute to energy consumption and increased carbon emissions. Organisations must adopt sustainable practices, like using energy-efficient algorithms or powering data centers with renewable energy. This reduces the environmental impact of AI while still benefiting from its capabilities.

Transparency

Transparency in AI means making users of AI solutions understand how it works. It helps ensure that users know how decisions are made and can trust the outcomes, especially in areas like healthcare, finance and law. AI transparency is key to ethical use because it allows people to see if the system is fair and unbiased. However, absolute transparency can also create risks like exposing the system to manipulation. At the same time, transparency must balance the need to protect system integrity against manipulation.

Bias and Discrimination in AI systems.

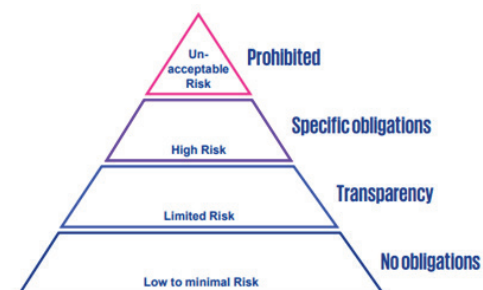
AI systems are designed to mimic human intelligence, as a result they can unintentionally inherit human biases. AI systems learn from data. They examine big data and identify patterns which help in making decisions or predictions. Therefore, if the data is biased, the AI will learn that bias and it will be reflected in its decisions. This is known as historical bias. Bias in AI can lead to discrimination. This means some people could be treated unfairly based on factors like race, gender or socioeconomic status. For example, it has become common for companies to use AI to screen job applications. If the AI is trained on data from previous recruitment practices that favoured men over women, it may unfairly reject qualified female candidates causing them to miss out on opportunities. Similarly, banks are also adopting the use of AI systems in loan approval processes. Again, if the AI is trained using data that is based towards a particular race or income, certain groups may be unfairly denied loans, even if they are financially capable of repaying them.

Bias in AI erodes public trust and can have damaging effects. The benefits of technology may be lost in sectors such as health care and finance where fairness and trust are invaluable. To prevent bias, organisations should ensure their AI systems use diverse datasets and regularly audit their systems for fairness. Ethical AI development is key to creating a future where technology serves all of society, without reinforcing the inequalities of the past.

Regulatory Frameworks

Governments and regulatory bodies worldwide are developing frameworks to ensure AI benefits society while minimising risks. Governments are dealing with balancing innovations in AI and its associated ethical responsibility.

On 1 August 2024, the European Artificial Intelligence Act officially came into effect. It marked a significant milestone in the European Union's effort in fostering responsible AI development and deployment in the EU. The Act aims to protect fundamental rights, democracy, the rule of law and environmental sustainability from high-risk AI, while boosting innovation and establishing Europe as a leader in the field. Furthermore, the Act provides developers and users with clear requirements and obligations regarding specific uses of AI while reducing administrative and financial burdens for businesses. The Act categorizes AI systems into four risk levels:



The act placed stricter regulations for higher-risk technologies such as biometric identification, systems used in medical diagnosis, autonomous surgical robots and systems that impact human rights.

In the USA, the White House released the Blueprint for an AI Bill of Rights in 2022 which addresses issues like transparency, discrimination, and data privacy. The AI Bill of Rights is made up of five principles to help guide the design, use, and deployment of automated systems, ensuring public rights protection in the age of artificial intelligence.

- The first principle of the Bill of Rights advocates that AI systems should be developed with input from the communities they impact to enhance safety and effectiveness.
- The next principle encourages companies to avoid algorithmic discrimination by ensuring that their AI systems treat everyone fairly.
- The third principle emphasizes data privacy, which means giving individuals greater control over their information and requiring clear consent for data collection.
- The fourth principle states that people should be notified when AI is involved in decisions that affect them and they should be provided with explanations of the role of AI.
- Finally, the fifth principle advocates for the provision of human alternatives which will allow individuals to opt out of AI systems in favour of human interaction when possible and encouraging feedback on their experiences.

In Ghana, the Minister of Communications and Digitalization highlighted the country's efforts to develop its AI regulatory framework. The Minister emphasized that AI at its core is not problematic, but its unregulated use could lead to ethical concerns. During the 75th edition of the Annual New Year School and Conference hosted by the University of Ghana in November 2023, the Minister announced that a draft AI policy had been formulated after significant stakeholder consultation and would soon be presented to Cabinet for approval. This policy aims to regulate AI for good, ensuring it serves the public interest by enhancing healthcare, transportation, and energy efficiency while addressing potential risks. As the Minister noted, "We are looking at AI for good," emphasizing the need to build regulatory structures that foster innovation while safeguarding societal interests.

Addressing Ethical Challenges Posed by AI

To address the ethical challenges and risks posed by AI systems, organisations can incorporate the following assessment and oversight processes into their AI governance and risk frameworks.

1. Perform AI readiness assessment - A high level review to discover the readiness of your organisation for implementing AI solutions at scale.
2. Establish a set of standards the organisation commits to – An AI ethics code of conduct, relevant to employees, customers and communities could be developed.
3. Establish ethics committees/boards to provide independent oversight, advice, assessment and monitoring of the ethics of AI systems throughout its lifespan.
4. Proactively engage relevant stakeholders throughout the AI lifecycle to ensure the system is addressing the right needs, in the right way.
5. Adopt a balanced and proportionate approach to ethical risk management and human oversight. The ethics of AI is contextual, because it is driven by cultural values, norms and the specific use of the system. As algorithms are designed to continuously learn from experience – the ethics of an AI system is also highly dynamic. This means that the ethical risks of AI need to be regularly and proactively monitored as well as subjected to targeted and proportionate oversight and due diligence.
6. Seek independent assurance of the ethics and broader trustworthiness of AI by putting in place regular reviews of the ethics of AI systems in your organisation.

The Future of Ethical AI

As AI continues to evolve, it is expected to become a bigger part of our everyday lives. AI is already making tasks easier, improving various sectors and transforming industries. In the future, ethical AI will become even more important. We can expect AI to be designed with fairness as more people become aware of the biases that can exist in AI. This means using more diverse data to train AI, testing systems for fairness and making sure that no group is left behind. This also means setting clear policies and laws that could require organisations to explain how their AI systems make decisions or ban the use of AI in ways that could cause disadvantages such as biased hiring practices or unfair loan approvals.

The future of ethical AI involves collaboration between humans and AI. Rather than replacing people, AI will assist in tasks that enhance productivity. For example, AI could handle repetitive tasks in manufacturing, allowing workers to focus on more complex responsibilities. The future also involves building trust in AI. This means giving people control over how AI interacts with their personal data and ensuring AI systems respect privacy. When people can trust AI, they will be more willing to embrace its benefits.

AI systems are continuously learning and evolving which means that new ethical challenges will emerge. As we move forward, our approach to ethics must also evolve. The future of ethical AI is about balance—ensuring that we keep pushing forward with innovation while protecting people’s rights, privacy and well-being.

Conclusion

AI has shown the ability to impact and improve multiple areas of life. However, its use gives rise to a myriad of ethical concerns. To ensure it benefits everyone, it should be developed and used responsibly. With the proper attention to ethical considerations, AI can be used for the benefit of all and not just a few. By recognising these issues, we can take steps to mitigate them. We have seen how bias and discrimination can sneak into AI systems, causing a socio-economic disadvantage to certain groups of people. Governments and organisations are beginning to establish regulations aimed at ensuring AI is transparent, accountable and safe for all. Our ability to unlock AI’s full potential will depend on how well we work together to uphold ethical standards. With the right approach, AI can shape a future driven by technology and guided by ethics.

How KPMG can help

KPMG understands that responsible and ethical AI involves complex business, regulatory and technical challenges and we are committed to helping clients put it into practice properly. We combine our deep industry experience, modern technical skills, leading solutions and robust partner ecosystem to help business leaders harness the power of AI in a trusted manner — from strategy and design through to implementation and ongoing operations. KPMG’s Trusted AI aims to assist organisations navigate the complexities of integrating AI technologies while

managing associated risks effectively. Trusted AI prioritises ethical considerations, addressing questions of privacy, consent and the responsible use of data. Trusted AI ensures that our clients are equipped to adapt and thrive within a landscape of AI ethical concerns, no matter what the future holds. We have assisted clients in diverse ways including performing:

- AI Risk assessment: We tailor frameworks designed to embed Trusted AI principles into every stage of AI initiatives through AI Readiness/ Maturity Assessments, AI Strategy Reviews or Trusted AI Strategy Implementation.
- AI Risk transformation: We help our clients to transform their governance and organisation in order to ensure accountability throughout the AI lifecycle as well as processes, controls and technologies to integrate Trusted AI into their end-to-end model management.
- AI Regulation and compliance: We perform comprehensive checks to align AI practices with international and industry-specific regulations as well as continuous monitoring and updates to keep our client’s AI systems compliant in a rapidly evolving regulatory landscape.
- AI Assurance: We test, examine evidence and report on management process, controls and claims regarding responsible use of AI technologies. Through our AI diagnostics reviews, we perform AI model control testing and attestation.

Wherever you are in your responsible AI journey, we can tailor our considerable experience, field-tested approach and innovative solutions to your unique needs and challenges, helping you to accelerate the value of AI with confidence.

Contact Us:



Andy Akoto
Partner, Advisory
KPMG in Ghana
E: aakoto@kpmg.com



Stephen Enebeli Annan
Manager, Forensics
KPMG in Ghana
E: stephenannan@kpmg.com



Michael Okai
Associate, Forensics
KPMG in Ghana
E: mokai@kpmg.com