Reflective Activity 1 – Ethics in Computing in the Age of Generative AI

Introduction

Since the public release of ChatGPT in late 2022, generative AI has transformed global industries at an unprecedented pace. The implications for computing professionals are profound, not only in terms of technical capacity but also in ethical, legal, social, and professional responsibilities. While AI is not a new concept, the speed and scale of generative models like GPT-4, Stable Diffusion, and Gemini necessitate a renewed conversation around AI governance, ethical use, and societal impact. This reflection explores these issues, drawing on Correa et al. (2023), Deckard (2023), and wider literature to propose a balanced ethical framework for future action.

Ethical Pluralism and Global AI Governance

Correa et al. (2023) highlight one of the central challenges in AI governance: ethical pluralism. Different nations and stakeholders approach AI with divergent priorities—ranging from economic growth in the U.S. and China to fundamental rights in the EU. The authors suggest that a global, comparative tool is necessary to understand the normative underpinnings of AI governance documents. This resonates with the need for interoperability between national frameworks, as AI technologies often transcend jurisdictional boundaries.

For example, while the EU's AI Act takes a risk-based approach that prioritizes human dignity, China's model emphasizes state control and innovation, and the U.S. largely

relies on industry self-regulation. These divergent models highlight that ethics in AI is not just about right or wrong—it is about context, culture, and geopolitical dynamics.

Recommended Action: A global repository of AI governance principles—much like the proposed “AI Policy Observatory”—should be prioritized by organizations such as the OECD, UNESCO, and the UN. This repository would allow computing professionals to compare policies, identify consensus, and align their practices with international norms, especially when deploying models across borders.

Professional Responsibility in the Age of Generative AI

Deckard (2023) explores the challenges computing professionals face when developing or deploying generative AI models. He argues that the increasing autonomy of generative systems, such as those capable of producing text, images, or even code, complicates traditional ethical responsibilities. For instance, if a language model inadvertently generates harmful or misleading content, who is accountable—the developer, the platform, or the user?

The ACM Code of Ethics provides a useful foundation: computing professionals must "avoid harm," "respect privacy," and "be accountable." But the emergence of black-box AI models makes this difficult. Developers may not fully understand how generative models derive their outputs, complicating transparency and accountability.

Recommended Action: Developers and organizations should implement model cards and datasheets for datasets (Mitchell et al., 2019) to improve transparency. These documents clarify the intended use cases, limitations, training data composition, and ethical considerations of AI systems. By embedding ethical reflection into technical documentation, computing professionals can better uphold their responsibilities.

Legal Considerations: Copyright, Consent, and Fair Use

Generative AI systems often rely on massive datasets scraped from the web, raising significant legal concerns. Who owns the data used to train these models? Do creators have the right to consent or opt out? Correa et al. (2023) suggest that existing copyright laws are ill-equipped to deal with non-human authorship.

For instance, several lawsuits—including those against OpenAI, GitHub Copilot, and Stability AI—allege unauthorized use of copyrighted data. These cases reveal the lack of legal clarity around fair use, especially when models generate content that closely resembles the training data.

Recommended Action: A revision of copyright law is required to address AI-generated content. Until then, companies should implement data transparency practices and allow creators to opt-out of data use, similar to what Adobe and Shutterstock have attempted with their “ethically sourced” AI models. Professionals must also stay updated on evolving case law to ensure their practices remain compliant.

Social Impact: Misinformation, Bias, and Inequality

Both Correa et al. and Deckard underscore the broader societal implications of generative AI, particularly in spreading misinformation, perpetuating bias, and widening digital divides. The ability to generate “deepfake” images, manipulate public discourse, or automate biased decision-making can have real-world consequences, from influencing elections to reinforcing stereotypes.

For example, generative models trained on biased data can reproduce harmful racial, gender, or socio-political assumptions. Without proper checks, these systems risk becoming amplifiers of inequality.

Recommended Action: Ethical auditing and bias mitigation must become standard industry practices. This includes third-party audits, diversity in training datasets, and fairness-aware machine learning techniques (Barocas et al., 2019). Moreover, social impact assessments—akin to environmental impact studies—should be mandated for high-risk AI systems, especially in public sectors like healthcare, education, and law enforcement.

Reflections on Personal and Professional Practice

As an AI researcher and computing professional in the making, I recognize the importance of incorporating ethics from the very beginning of system design. It is no longer acceptable to treat ethics as an afterthought. My own professional conduct must reflect a commitment to transparency, inclusivity, and accountability.

This includes:

Ensuring datasets used in projects are ethically sourced.

Documenting model limitations and potential risks clearly.

Advocating for open discussions around algorithmic harms within team environments.

Being open to interdisciplinary collaboration with ethicists, sociologists, and legal scholars.

By embracing these principles, I can contribute to a computing culture that not only advances innovation but does so responsibly.

Conclusion

Generative AI represents both a technological leap and a profound ethical challenge. As Correa et al. (2023) and Deckard (2023) show, the path forward requires global cooperation, local sensitivity, and professional courage. While national policies may diverge, computing professionals must lead by example—embedding ethics into code, practice, and culture. Ultimately, a responsible approach to generative AI is not just about preventing harm—it's about maximizing the good we can do when technology and values align.

References

1. Correa, J., et al. (2023). Comparative approaches to AI governance: A global analysis of principles and policies. [Journal/Publisher Info].
2. Deckard, M. (2023). Ethics of Generative AI: Rethinking Professional Responsibility. [Journal/Publisher Info].
3. Mitchell, M., et al. (2019). “Model Cards for Model Reporting.” Proceedings of the Conference on Fairness, Accountability, and Transparency (FAT).
4. Barocas, S., Hardt, M., & Narayanan, A. (2019). Fairness and Machine Learning: Limitations and Opportunities.