



Delft University of Technology

Generative AI in Software Engineering Must Be Human-Centered The Copenhagen Manifesto

Russo, Daniel; Baltés, Sebastian; van Berkel, Niels; Avgeriou, Paris; Calefato, Fabio; Cabrero-Daniel, Beatriz; Catolino, Gemma; Izadi, Maliheh; Vasilescu, Bogdan; More Authors

DOI

[10.1016/j.jss.2024.112115](https://doi.org/10.1016/j.jss.2024.112115)

Publication date

2024

Document Version

Final published version

Published in

Journal of Systems and Software

Citation (APA)

Russo, D., Baltés, S., van Berkel, N., Avgeriou, P., Calefato, F., Cabrero-Daniel, B., Catolino, G., Izadi, M., Vasilescu, B., & More Authors (2024). Generative AI in Software Engineering Must Be Human-Centered: The Copenhagen Manifesto. *Journal of Systems and Software*, 216, Article 112115. <https://doi.org/10.1016/j.jss.2024.112115>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy


Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.



Generative AI in Software Engineering Must Be Human-Centered: The Copenhagen Manifesto

Daniel Russo^a, Sebastian Baltes^b, Niels van Berkel^a, Paris Avgeriou^c, Fabio Calefato^d, Beatriz Cabrero-Daniel^e, Gemma Catolino^f, Jürgen Cito^g, Neil Ernst^h, Thomas Fritzⁱ, Hideaki Hata^j, Reid Holmes^k, Maliheh Izadi^l, Foutse Khomh^m, Mikkel Baun Kjærgaardⁿ, Grischa Liebel^o, Alberto Lluch Lafuente^p, Stefano Lambiase^g, Walid Maalej^q, Gail Murphy^k, Nils Brede Moe^r, Gabrielle O'Brien^s, Elda Paja^t, Mauro Pezzè^u, John Stoubly Persson^a, Rafael Prikladnicki^v, Paul Ralph^w, Martin Robillard^x, Thiago Rocha Silvaⁿ, Klaas-Jan Stol^y, Margaret-Anne Storey^h, Viktoria Stray^z, Paolo Tell^t, Christoph Treude^{aa}, Bogdan Vasilescu^{ab}

^a Aalborg University, Denmark

^b University of Bayreuth, Germany

^c University of Groningen, The Netherlands

^d University of Bari, Italy

^e University of Gothenburg & Chalmers, Sweden

^f University of Salerno, Italy

^g TU Wien, Austria

^h University of Victoria, Canada

ⁱ University of Zurich, Switzerland

^j Shinshu University, Japan

^k University of British Columbia, Canada

^l Delft University of Technology, The Netherlands

^m Polytechnique Montreal, Canada

ⁿ University of Southern Denmark, Denmark

^o Reykjavik University, Iceland

^p Technical University of Denmark, Denmark

^q University of Hamburg, Germany

^r SINTEF, Norway

^s University of Michigan, USA

^t IT University of Copenhagen, Denmark

^u USI Università della Svizzera Italiana, Switzerland

^v Pontifícia Universidade Católica do Rio Grande do Sul, Brazil

^w Dalhousie University, Canada

^x McGill University, Canada

^y University College Cork, Ireland

^z University of Oslo, Norway

^{aa} Singapore Management University, Singapore

^{ab} Carnegie Mellon University, USA

Preamble. The advent of Generative Artificial Intelligence—systems that can produce human-like content such as text, music, visual art, or source code—marks not only a significant leap for Artificial Intelligence (AI) but also a pivotal moment for software practitioners and researchers. The role of software engineering researchers and practitioners in adopting the technologies that shape our world is critical. Historically, the human aspects of developing software have been treated as secondary to more technical innovations. However, the emergence of Generative AI will simultaneously enhance human capabilities while

surfacing complex ethical, social, legal, and technical challenges.

While primarily aimed at software engineering (SE) researchers and practitioners, who are at the forefront of integrating Generative AI into our digital infrastructure, this manifesto equally underscores the ripple effects such innovations have on end-users of software, shaping their experiences and interactions in profound ways.

SE practitioners and researchers have a moral duty to help the world navigate this enduring landscape of new AI technologies ethically and sustainably, and foster values such as fairness, transparency, societal

E-mail address: daniel.russo@cs.aau.dk (D. Russo).

<https://doi.org/10.1016/j.jss.2024.112115>

Available online 31 May 2024

0164-1212/© 2024 Elsevier Inc. All rights reserved, including those for text and data mining, AI training, and similar technologies.

wellbeing, and environmental resilience. This transformative era challenges us to extend our focus beyond technical expertise, integrating human values and ethical considerations into the fabric of our technological advancements. This is in addition to maintaining our existing obligations, such as adherence to codes of conduct (e.g., ACM Code of Ethics and Professional Conduct¹ or the IEEE Code of Ethics²) and compliance with legal frameworks (e.g., EU AI Act³).

By establishing this manifesto, we seek to catalyze a shift in how Generative AI is conceived, developed, and applied within SE—a shift that reaffirms the primacy of human dignity, agency, and collective wellbeing in the face of rapid technological change.

Core Values of the Copenhagen Manifesto

- 1. Responsibility and Ethics:** Emphasizing the duty to develop Generative AI in SE in a manner that does not bring (physical, emotional, or financial) harm to living beings.
- 2. Human-Centricity and Agency:** Prioritizing human needs and autonomy when designing and using Generative AI in SE.
- 3. Transparency and Equity:** Advocating for transparent, understandable, reproducible, and verifiable Generative AI in SE, including communications about research in this area, and equitable access and impact.
- 4. Inclusivity and Continuous Learning:** Adopting, adapting, and empirically researching Generative AI in SE must inherently foster a culture of diversity, inclusivity, and continuous learning, embracing diverse, multidisciplinary perspectives for software engineering's advancement.
- 5. Environmental Sustainability:** Committing to the responsible selection and use of Generative AI technologies, with a focus on minimizing environmental impact through energy efficiency and sustainable practices.

Bridging our core values to actionable principles, the Copenhagen Manifesto presents a roadmap for embedding these commitments into SE practices with Generative AI. These principles serve as concrete steps towards realizing a human-centered approach, ensuring that technological advancement is not only innovative but also equitable, transparent, and socially and environmentally responsible. As we move forward, these guidelines will steer us in applying Generative AI within SE in a manner that enhances the human experience, fostering an environment where technology truly serves society.

Principles

- 1. Responsible Management of Generative AI in SE:** Actively engage in responsibly managing and continuously evaluating Generative AI technologies, ensuring their alignment with ethical standards and societal needs.
- 2. Human Sovereignty Over AI:** Ensure that Generative AI technologies enhance, rather than replace, human decision-making and creativity, prioritizing human oversight in Generative AI development and application.
- 3. Two Sides of Generative AI:** Assess both the benefits and harms of Generative AI, promoting its use with caution and responsibility to avoid unintended consequences.
- 4. Sociotechnical Responsibility:** Integrate social and technical considerations in the development of Generative AI-powered applications, aiming for solutions that are beneficial and respectful to all stakeholders.
- 5. Transparency and Fairness:** Implement transparent processes that actively identify and mitigate Generative AI-related biases,

ensuring fairness, accountability, and trustworthiness in Generative AI applications.

- 6. Sustainability and Environmental Impact:** Select and advocate for Generative AI models and practices known for their lower environmental impact, emphasizing long-term sustainability.
- 7. Equitable Impact of Generative AI in Software:** Champion Generative AI that fosters inclusivity and equity, working to eliminate disparities and support empowerment across diverse communities.
- 8. Ethical Engineering and Education:** Uphold ethical standards in both engineering practices and education, ensuring future generations of engineers are prepared to make ethical decisions in Generative AI use.
- 9. Complex Adaptive Nature of Generative AI:** Recognize and plan the evolving nature of Generative AI systems by committing to ongoing learning and adaptation in their design and application, ensuring they remain aligned with ethical and societal values.
- 10. Empirical Evaluation of Professional Practices:** Support and conduct empirical research to base Generative AI adoption and practices on evidence, enhancing the positive impact of Generative AI on society and addressing any challenges.
- 11. Enhancement and Wellbeing through Adaptive Generative AI in SE:** Develop Generative AI that amplifies the capabilities of software engineers, promoting professional growth and wellbeing.
- 12. Public Awareness and Discourse:** Engage in and foster public discussions about Generative AI, aiming to educate and inform while combating misinformation, thereby shaping a balanced and factual narrative around Generative AI in SE.

As we explore the possibilities of Generative AI in SE, we must keep our eyes open not only to its incredible opportunities but also to its considerable risks.

Generative AI is reshaping the landscape, affecting how we live and work. Our mission is to steer Generative AI in SE towards fairness, transparency, and justice, ensuring a future where technology lifts us up together, rather than divides us.

Talking about responsibility, ethics, transparency, fairness, and focusing on everyone's wellbeing means insisting that all technology, including Generative AI, serves the common good. It is about developing and applying Generative AI with respect for each person's rights and dignity to build a world that is better for our collective effort.

In conclusion, this manifesto represents our collective commitment to a future where technology is guided by the values and principles we have laid out. Also, it is a pledge to advance software engineering not just as a field of technical expertise but as a domain profoundly connected to ethical integrity, social responsibility, and human wellbeing.

Declaration of competing interest

None.

Acknowledgments

Our sincere appreciation goes to the Alfred P. Sloan Foundation and the Carlsberg Foundation for their invaluable support of the Copenhagen Symposium on Human-Centered Software Engineering AI (Grants G-2023-21020 and CF23-0208 to Daniel Russo). Held at Aalborg University – Copenhagen from November 26-29, 2023, this symposium was the birthplace of this manifesto, and contributions from these foundations have been crucial in shaping its vision and direction.

¹ <https://www.acm.org/code-of-ethics>.

² <https://www.ieee.org/about/corporate/governance/p7-8.html>.

³ <https://data.consilium.europa.eu/doc/document/ST-5662-2024-INIT/en/pdf>.