


AI Policy Research Summit, Stockholm,
November 2024

Roadmap for AI Policy Research



Introduction

With Artificial intelligence (AI) development and adoption advancing at an unprecedented pace, policymakers and regulators are encountering both significant challenges and opportunities. The challenges emerge from the disconnect between the fragmented & at times siloed policy & technology landscapes, whilst the opportunities relate to novel insights and capabilities afforded to decision-makers to strengthen the evidence base for sustainable policy. While AI offers transformative potential, it also poses substantial risks, such as biases, inequalities, and threats to privacy and security. In this context, AI policy research has emerged as an essential guide to navigating the complex interplay between technological innovation and societal impact. It ensures that advancements in AI align with ethical, legal, and social priorities. AI policy research provides the evidence base needed to address these challenges, fostering accountability, transparency, fairness, and inclusivity in AI governance. It also helps anticipate future regulatory needs, bridge the gap between stakeholders, and ensure that AI technologies are deployed responsibly and equitably, contributing to sustainable development and the public good.

This roadmap, developed through collaborative discussions at the recent [AI Policy Research Summit](#), reflects a shared vision for advancing research on responsible AI policy and governance. It emphasizes the critical role of policy research in ensuring that AI development is guided by robust evidence, ethical considerations, and a commitment to sustainability and inclusivity. By prioritizing transparency, accountability, and the well-being of humans and the planet, this roadmap highlights how research can inform global approaches to AI governance, addressing complex societal needs and ethical challenges. It serves as a guiding framework for stakeholders across academia, industry, government, and civil society to collaborate in generating actionable insights and evidence-based strategies, noting that while evidence-based AI policy often draws on data-driven research, it equally values critical theoretical insights and fundamental rights approaches, ensuring a holistic understanding that extends beyond the purely quantifiable.



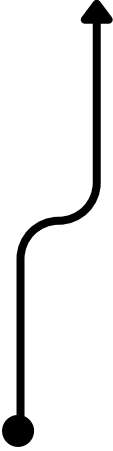
Core Principles

Guiding principles for AI policy research should ensure rigorous adherence to scientific and disciplinary methodologies. Moreover, these principles aim to produce policy recommendations that not only reflect a nuanced understanding of AI's potential but also emphasize the importance of balancing transformative benefits with responsible and equitable implementation. By doing so, they provide a framework that supports the ethical and effective integration of AI into society. Key principles leading AI policy research include:

Human and Planetary Welfare: Research must prioritize the development of policies that enhance individual and collective well-being while promoting [human rights](#) and environmental sustainability. This includes systematically addressing direct and indirect environmental impacts of AI systems, such as [biodiversity](#), [the ecosystem](#), energy and water consumption, and electronic waste, and aligning AI deployment with the United Nations' Sustainable Development Goals (SDGs). In particular, AI policy research should address issues of exploitation within AI development. This concerns both the hardware (mineral extraction) as well as exploitative work practices where data workers often have little to no autonomy and where psychologically and physically demanding tasks, e.g., removing toxic content, are outsourced without proper compensation and labour rights.

Accountability and Transparency: AI policy research should advocate for, and provide the tools and methods to support, decision-making processes that are transparent, evidence-based, and accountable. This entails creating frameworks for accessible, relevant, and reliable explainability in AI systems and ensuring that policymakers and stakeholders understand the risks, assumptions, and trade-offs involved in AI governance. The transparency principle underscores the importance of traceability in AI system design, implementation, deployment, and, whenever possible, reproducibility.

Inclusivity, Diversity, and Capacity Building: Effective AI policy research must address disparities in access to AI technologies, their benefits, and AI-related decision-making processes especially in underrepresented or historically marginalized communities. This includes examining issues such as data bias, unequal distribution of AI advancements and resources or AI risks, and systemic barriers to participation in AI ecosystems on local, regional and global levels. Promoting diversity in AI governance is essential to ensuring that AI serves a wide range of societal interests.



Ethical Research Practices and Institutional Responsibility: AI policy research and collaborations must adhere to rigorous ethical standards that ensure innovation aligns with societal well-being. This involves upholding principles of fairness, institutional responsibility, accountability, human-centered values, human rights, robustness, safety, and inclusiveness throughout the research process. Researchers should be guided by established frameworks, such as those from OECD and UNESCO, to maintain transparency, integrity, and methodological rigour. By fostering trust and credibility, these practices ensure not only responsible research but also institutional accountability and a commitment to ethical governance in shaping the future of AI.

Ethical Governance: Governance should be guided by ethical considerations that balance innovation with societal impact, following principles of fairness, accountability, human-centered values, human rights, robustness, safety, and inclusiveness as advocated by principles and guidelines such as OECD and UNESCO. AI policy research provides the much-needed scientific evidence to support and refine these guidelines. By grounding governance principles in scientific theories and empirical research, AI policy research ensures that ethical frameworks remain actionable, measurable, and responsive to evolving challenges, fostering responsible AI policy and governance.

Equitable economic growth: AI policy research should focus on ensuring that the economic benefits of AI are distributed fairly across society, promoting opportunities for underrepresented communities and fostering inclusive growth. This includes addressing systemic economic inequalities exacerbated by AI adoption, supporting policies that create equitable access to AI-driven markets and resources, and encouraging sustainable market practices that prevent monopolistic dependencies. By prioritizing fairness and inclusivity in economic frameworks, AI policy can contribute to shared prosperity and long-term economic resilience.

Priorities for Impactful AI Policy Research

Addressing pressing challenges in AI policy requires focused and collaborative research efforts to understand transformative benefits and ensure AI governance aligns with societal needs. These efforts must bridge gaps between stakeholders, anticipate future developments, and foster cross-border cooperation to guide responsible and inclusive AI innovation. In particular, we identify the need for research on:

Transboundary AI Governance: Effective AI governance requires exploring how principles such as transparency, accountability, and trust can be interpreted, understood, and implemented across borders, in particular in regional harmonised initiatives, without assuming universal uniformity. Research should focus on identifying flexible mechanisms for harmonising standards among diverse actors in the AI ecosystem. This includes examining context-sensitive approaches to governance that respect local values and priorities while fostering international cooperation. By exploring horizontal governance strategies, studies can clarify operational norms and facilitate collaboration, even in the absence of globally binding legislation.

Defining and Measuring the Benefits of AI: As AI continues to shape society, how do we define and measure its benefits? AI promises to enhance well-being, efficiency, and innovation, but who truly gains, and who is left behind? Are different benefits in conflict, and can they be meaningfully qualified or quantified? Answering these questions requires reflection on what it means for AI to be beneficial and how we assess its true impact on individuals and communities.

Foresight and Proactive Regulation: AI policy research must anticipate technological advancements and their implications to guide regulatory frameworks proactively and leverage the use of AI to identify and monitor signals for positive change. This includes conducting foresight exercises to predict emerging trends and their societal impacts, enabling policymakers to prepare adaptable and forward-thinking regulations. It also includes understanding the impact of AI narratives on public attitudes towards AI and perception of humanity's value and role in society.

Codes of Conduct: AI policy research can also inform the development of codes of practice for different types of stakeholders, including researchers, developers, deployers and policymakers, ensuring that governance evolves alongside innovation, rather than in reaction to it.

Stakeholder Collaboration: Research in AI policy can identify best practices for stakeholder engagement, develop frameworks for collaborative dialogue, and create scenarios that align diverse interests. By bridging gaps between stakeholders, research ensures that governance decisions are informed by comprehensive, multi-perspective insights.

Sectoral Focus: Research can prioritize areas of high impact for individuals, communities, and regions, such as sustainability, health, job access, or security, ensuring that an AI solution is not taken for granted but part of a critical and participatory process (that is, addressing "question zero"). Comparative analyses of regulatory approaches from fields such as environmental policy and healthcare governance can provide valuable insights.

Guiding Actions and Roadmap

To develop a comprehensive, collaborative, and cross-border effort towards advancing AI policy research, we propose the following immediate and long-term actions:

- 1** **Establishing a Community of Practice:** Develop networks and platforms for researchers, policymakers, and industry leaders to share best practices, foster collaboration, and address challenges in AI policy research.
- 2** **Visiting AI Policy Fellowships:** Create fellowship opportunities to build stronger connections between researchers and policymakers, enabling practical exchange of knowledge and alignment of research with policy needs.
- 3** **Annual AI Policy Summit:** Host an annual forum to bring together stakeholders for discussions, knowledge-sharing, and collaborative reflection on AI policy research progress and priorities.
- 4** **Call to Action on AI Policy:** Engage a broad range of stakeholders through targeted campaigns to promote the adoption of ethical, evidence-based AI policies and governance frameworks.
- 5** **AI Policy Briefs:** Develop and distribute concise, actionable, and timely briefs to communicate research findings and policy recommendations effectively to decision-makers.
- 6** **Student and Staff Exchange and Capacity Building:** Facilitate academic exchanges, exchanges between industry and academia, and interdisciplinary programs to train future leaders in AI policy research and governance, fostering a globally connected and informed community.
- 7** **AI Literacy:** Promote AI literacy for people of all ages through accessible media, educational programs, and public engagement efforts, ensuring a widespread understanding of AI's impact, opportunities, and risks.

Commitments to Responsible Research and Governance

We commit to advancing AI policy research that meets the growing need for responsible governance, grounded in ethical, transparent, and evidence-based practices to shape inclusive and trustworthy policies, including:

- Responsible Research:
 - Transparency in methods, funding, and stakeholders.
 - Rigorous ethical standards and research integrity.
 - Advocacy for AI policies that avoid undue industry influence.
- Responsible Governance:
 - Clear communication with policymakers to build trust.
 - Avoidance of undue influence from commercial entities.
 - Evidence-based consultation with relevant stakeholders to shape inclusive policies.

Call to Action

The roadmap calls for unity among stakeholders to strengthen AI policy research efforts. By engaging diverse groups, including policymakers, researchers, civil society, and industry, this initiative seeks to ensure that AI advancements serve the public good, promoting accountability, sustainability, and inclusivity.

We invite your feedback to help refine these strategies, ensuring they effectively support meaningful contributions to the global dialogue on AI policy and governance, and shaping the role of AI policy research in addressing the challenges and opportunities in this evolving field.

Authors

(Authors in alphabetical order after the two leading authors)

- Virginia Dignum, Director AI Policy Lab, Umeå University
- Catherine Régis, Professor of Law, University of Montreal, Canada - CIFAR AI Chair, Mila
- Kerstin Bach, Director NorwAI Research Center
- Stefan Buijsman, Associate Professor of Ethics of Technology, Delft University of Technology
- André P. L. F. de Carvalho, Center for Artificial Intelligence and Machine Learning (CIAAM) and Institute of Mathematics and Computer Science (ICMC), University of São Paulo, Brazil
- Ginevra Castellano, Professor, Director of the Uppsala Social Robotics Lab, Department of Information Technology, Uppsala University
- Frank Dignum, Director of the Center for Transdisciplinary AI Research for Good (TAIGA), Umeå University
- Elizabeth Farries, Director, Centre for Digital Policy, University College Dublin
- Fosca Giannotti, Scuola Normale Superiore, Pisa, Italy
- The Anh Han, Professor of Computer Science, Lead of Center for Digital Innovation, Teesside University
- Natali Helberger, Distinguished University Professor Law & Digital Technology, University of Amsterdam
- Isadora Hellegren, Senior Project Manager, Research, Mila - Quebec AI Institute
- Geert-Jan Houben, TU Delft
- Anna Jahn, Senior Director, Public Policy and Inclusion, Mila - Quebec AI Institute
- Somya Joshi, Associate Professor Stockholm University & Head of Division: Global Agendas, Climate & Systems, Stockholm Environment Institute, Sweden.
- Mamadou Lamine Sarr, ALIVE Lab, Cheikh Hamidou Kane Digital University (UNCHK), Senegal
- Dave Lewis, Professor in Computer Science, ADAPT Centre at Trinity College Dublin, Ireland



Authors

(continuation)

- Anna-Sara Lind, Uppsala University and WASP-HS
- Neema Lugangira, Chair of African Parliamentary Network on Internet Governance (APNIG) and Member of Parliament Tanzania
- Helena Gorete Silva Moniz, University of Lisbon, Center for Responsible AI
- Laurens Naudts, AI, Media & Democracy Lab – Institute for Information Law, University of Amsterdam
- Seydina Moussa Ndiaye, Director AI Living Lab ALIVE, Cheikh Hamidou Kane University, Senegal
- Julia Palma, CeADAR - Ireland's Centre for Artificial Intelligence, University College Dublin
- Dino Pedreschi, Università di Pisa, Italy
- Nitin Sawhney, Professor of Practice, Aalto University and Finnish Center for Artificial Intelligence (FCAI)
- Sara Spaargaren, Manager AI, Media & Democracy Lab, University of Amsterdam
- Nicole Tong, Stanford University
- Edward Tsoi, Co-Founder, AI Safety Asia
- Jason Tucker, Researcher at The Institute for Futures Studies. Visiting Research Fellow, AI & Society, Lund University and Affiliate, AI Policy Lab, Umeå University.
- Ysaline Bourguine de Meder, Head of Strategy, Effective Altruism Sweden

Design

Tatyana Sarayeva, Project Coordinator AI Policy Lab, Umeå University