
DIVE INTO AI: ETHICS

WHAT WE HEARD

Building Bridges

May 2025

TABLE OF CONTENTS

LAND ACKNOWLEDGEMENT

INTRODUCTION

SSUMMARY

SHARED THEMES

DISCUSSION

SPEAKER BIOGRAPHIES

BUILDING BRIDGES

LAND ACKNOWLEDGEMENT

The workshop series began with acknowledging and respecting the First Nations, Metis, and Inuit peoples on whose traditional territories the Canadian Integrated Ocean Observing System's (CIOOS) work is conducted across Canada.

We strive for respectful relationships with all Indigenous peoples as we search for collective healing, true reconciliation, and honouring these beautiful lands and seas together.

The virtual session hosts from CIOOS Atlantic were located in Kijipuktuk, otherwise known as Halifax, the traditional and unceded territory of the Mi'kmaq.

INTRODUCTION

This workshop offered an exploration of responsible AI, emphasizing the importance of fairness, transparency, accountability, and continuous oversight in both the development and deployment of AI systems. Presenters emphasized that bias, legal, and ethical obligations, along with human-centered values, must guide the use of AI. The discussion also underscored that AI is not a one-size-fits-all solution, and its risks and benefits must be weighed in each situation. The participants raised thoughtful questions about sustainability, governance, and preserving human agency, which point to a deep concern about how AI intersects with social and environmental systems. The presentations and discussion highlighted how building trustworthy AI requires ongoing collaboration, critical reflection, and a strong commitment to the public good.

The 'Dive Into AI: Ethics' workshop focused on disentangling the complex web of issues and considerations to navigate AI safely, responsibly, and ethically. This 2-hour virtual session featured three keynote speakers, followed by a panel discussion.

SUMMARY

Anne Nguyen, Chief AI Officer (Conseil de l'innovation du Québec)

Anne Nguyen's presentation focused on the importance of developing and using AI responsibly and ethically, especially for the benefit of future generations. She outlined the role of the Conseil in supporting innovation in Quebec, and shared that the province was asked by the government to lead a public conversation about responsible AI, which resulted in a significant report with practical recommendations. Anne highlighted that, while AI is a powerful tool that is capable of increasing efficiency and providing benefits in sectors such as healthcare, it is incumbent on people to decide what constitutes an acceptable use. Anne reinforced the idea that AI doesn't make decisions; it makes predictions based on data. Her presentation also touched on the history of AI, the strengths and limits of the technology, and the need for the human values of creativity, intention, and fairness to guide its use.

Building trustworthy and effective AI in government requires public engagement, policy alignment, and governance structures that prioritize transparency, accountability, and equity. Several themes emerged throughout the presentation, including: responsible and ethical AI use, human-AI collaboration, public awareness and education; Innovation as a collective effort; and AI as both opportunity and risk.

Jim Provost, Chief AI Officer (ScarlettNova)

Jim Provost's presentation, about the ethical challenges and risks of AI, focused on how training data may allow bias and discrimination to creep into AI systems. Because AI is built from large amounts of historical data and "known right answers," the quality and representativeness of said data are incredibly important. Using real-life examples, he demonstrated how biased data can lead to unfair or discriminatory outcomes. Jim emphasized that AI doesn't think for itself; it is advanced statistics predicting outcomes based on past patterns. He walked through methods to measure errors and bias in AI

systems, highlighting that fairness and quality must be actively measured, not just assumed. He also stressed that AI is not a “one-and-done” project, but an ongoing process that requires regular updates and evaluation. Throughout the talk, Jim encouraged thoughtful data sampling, legal awareness of discrimination, and being transparent about assumptions.

Ethical AI starts with humans taking responsibility for how data is collected, interpreted, and used. Several themes emerged throughout the presentation: bias and fairness in AI; the importance of measurement and evaluation; legal and ethical context; AI as a continuous product, not a one-time project; limitations of “clean” or “unbiased” data.

Frank Rudzicz, Associate Professor of Computer Science (Dalhousie University)

Frank Rudzicz called on his experience as an AI researcher and health-tech entrepreneur to explore the practical and ethical complexities of responsible AI, particularly in healthcare. He outlined the principles of transparency, fairness, accountability, and privacy, illustrating how real-world AI systems can unintentionally reinforce societal biases or compromise privacy, even when designed with good intentions. The need for rigorous data handling, robust fairness metrics, and awareness of unintended consequences was emphasized, and the tension between privacy protections and model performance for marginalized groups was highlighted, making a case for nuanced, context-sensitive approaches. He concluded by advocating for Canadian leadership in responsible AI and invited collaboration through a new AI institute at Dalhousie University dedicated to research, education, and capacity building.

Building responsible AI requires recognizing and managing the complex trade-offs between fairness, privacy, transparency, and real-world effectiveness. Several themes emerged throughout the presentation: core pillars of responsible AI; bias in AI and healthcare; privacy and fairness trade-offs; limits of generative AI hype; real-world applications of AI in health; ethical frameworks and governance for AI; and, building a responsible AI ecosystem in Canada.

SHARED THEMES

Bias and Fairness in AI

Bias is a critical issue in AI systems. Anne Nguyen and Frank Rudzicz addressed how bias arises from training data and systemic inequality, while Jim Provost emphasized fairness as a human value, a moral imperative, and a technical challenge that requires intentional design and oversight.

Responsible AI as a Guiding Principle

Each speaker emphasized that AI development must be governed by human values such as fairness, transparency, accountability, and privacy. Anne Nguyen focused on legal contexts and metrics for evaluating responsibility, Jim Provost framed responsibility in terms of public trust and human oversight, and Frank Rudzicz presented responsible AI use based on core ethical pillars and real-world governance frameworks.

Requirement for Ongoing Oversight and Iteration

The concept of AI requiring constant attention—that it is not a "build it and forget it" technology—runs through all three presentations. Anne Nguyen described AI as a continuous product needing evaluation and maintenance. Frank Rudzicz emphasized model monitoring, retraining, and the infrastructure required to support this, while Jim Provost tied the need for long-term responsibility to public education and ethical use.

DISCUSSION

Following the presentations was an open dialogue with our experts, during which the audience was invited to ask questions. The theme of responsible governance and accountability ran throughout, with audience questions reflecting concern for how AI impacts people and society, and highlighting a collective concern to ensure AI serves the public good, not just technological advancement. The following topics were addressed:

Balance the AI benefits with environmental cost

The environmental impact of AI, particularly large language models, is significant and growing, with projections suggesting that AI could consume as much energy as a small country by 2027. A core recommendation to address this is proportionality: matching the complexity of the tool to the task. Rather than using an advanced AI model to perform simple web searches or trivial tasks, be intentional and selective. Deploy AI only where the benefits truly justify the energy cost. Yet it is difficult to calculate an accurate cost-benefit analysis, as the current systems for recording environmental externalities associated with AI are inadequate to tally their true costs.

Maintaining control over data management and integrating ethical practices

Maintaining ethical data practices requires both structural and human considerations. In some cases, self-imposed organizational requirements automatically help enforce good practices. But ethical integration requires more than written rules; success often lies in the knowledge and capabilities of people. Implementation depends heavily on whether the team understands the technical tools available, how they interact, and how to apply them responsibly. Continuous education and collaboration on emerging standards are needed to ensure that ethical goals are operationalized.

DISCUSSION CONTINUED

Ethical considerations when sharing models trained on non-public Indigenous data

Non-public data from Indigenous communities raises both legal and cultural considerations that cannot be generalized or assumed, and often require consultation with experts in Indigenous law, data governance, and ethics. The decision to share models trained on such data must be guided by appropriate community consultation, respect for data sovereignty, and compliance with any applicable legal frameworks.

AI Use that does not Degrade Essential Human Skills

To mitigate skill loss, AI users must be aware of how and why they use AI. Even though AI may be capable of performing a task, there may be value in doing it manually, especially for tasks that exercise critical thinking or communication skills. Balancing automation with intentional practice of human skills is key to maintaining long-term capability and confidence in areas like writing, reasoning, and problem-solving.

Ethical Documentation of Methods during AI Development

Ethical documentation in AI development is more than simply tracking performance metrics; it requires recording the full process, from start to finish. Development tracking platforms and tools, such as MLflow and Weights & Biases, can assist in logging these decisions—including the models that were considered, the data used, and the performance metrics that were evaluated. Ethical documentation also requires that fairness-related metrics be recorded and that there is transparency about any limitations or tradeoffs. By embedding documentation in workflows and maintaining proper design practices, organizations can ensure traceability and accountability in how AI systems are built and deployed.

Improving Transparency and Explainability in AI Systems

Improving transparency in black-box AI systems requires the use of explainability tools. Methods such as SHAP, LIME, and

feature attribution techniques help illuminate the internal logic of models by allowing developers to determine the inputs that contribute to specific outputs. But these tools are imperfect and can sometimes give misleading results, and as such, their interpretations must be carefully managed.

Defining Responsible AI

Elected governments and public institutions should establish definitions of responsible AI through a democratic process. At the same time, both public and private organizations can create their own principles and standards, tailored to their values and objectives. Responsible AI entails setting clear ethical guidelines and governance structures. International bodies, academic committees, and standards organizations—such as the European Union—play a crucial role in shaping the discourse on responsible AI. While this conversation continues to evolve, disagreements about what constitutes fairness, bias, or responsibility are likely to persist. Therefore, institutions must be transparent about their definitions, and the public must remain actively engaged in shaping these definitions through input on policy and oversight.

Academia's Response to a Quickly Changing AI Landscape Dominated by Industry

It is becoming increasingly difficult to keep up with the rapid pace and scale of industrial AI development. Major technology companies have considerably more resources than educational institutions, and are better-equipped to manage the substantial cost of generative AI development. As a result, research has shifted from next-generation AI development to solving real-world problems with existing tools. Researchers are focusing on deployment challenges, fairness, explainability, and real-world applications in areas such as the environment and health. Universities are now playing an important role in validation. If you are working with AI, especially in oceans, you are likely working with a strong multi-million parameter machine learning model that does not require special computing capacity.

Keeping Up to Date with Rapid Changes in AI

Staying up to date on AI developments requires constant learning. Outside of formal channels such as academic papers and conferences, informal channels—including conversations with students or social media jokes—may serve as cues to explore emerging tools or ideas.

SPEAKER BIOGRAPHIES

Anne Nguyen, Chief AI Officer (Conseil de l'innovation du Quebec)

Anne is a champion for innovation, guided by the values of inclusion, equality, and humanity, whose career has been marked by numerous challenges and as many achievements. With a diploma from Harvard Business School in Business Analytics, Anne Nguyen has been in the role of Chief AI Officer with the Conseil de l'innovation de Québec since January 2024.

Jim Provost, Chief AI Officer (ScarlettNova)

Jim has over two decades of experience in software, automation, and data science. He received his B.Eng from McMaster and a Master's from Queen's University. He is a strategic adviser, a previous TEDx speaker, and a father of two.

Frank Rudzicz, Associate Professor of Computer Science (Dalhousie University)

Frank's work is in machine learning and healthcare, with a focus on natural language processing, speech recognition, and safe AI. His research has appeared in popular media such as Scientific American, Wired, CBC, the New York Times, and Nature.

ABOUT BUILDING BRIDGES

The Dive Into AI Workshops, hosted by CIOOS, are part of the Building Bridges project, announced in June 2024 and funded by Canada's Ocean Supercluster.



The Building Bridges project focuses on equipping organizations in the ocean sector with the tools and knowledge needed to implement AI solutions effectively. Building Bridges plans to lower AI adoption barriers through addressing multiple components of AI development.

Contact us at info@cioos.ca if you have any questions about Building Bridges or the Workshop sessions.

CIOOS
CANADIAN INTEGRATED OCEAN
OBSERVING SYSTEM



SIOOC
SYSTÈME INTÉGRÉ D'OBSERVATION
DES OcéANS DU CANADA

Email
info@cioos.ca

Website
www.cioos.ca