


CAMBRIDGE JOURNAL OF ARTIFICIAL INTELLIGENCE



Vol. 1 Issue 3
July 2026

© *Cambridge Journal of Artificial Intelligence*.

All rights reserved.

Published by *Cambridge Journal of Artificial Intelligence*, Cambridge, United Kingdom.

The cover artwork is *Corruption 1* by Kathryn Conrad. The image represents both the way that AI models training on their own output leads to model collapse and the way that corrupted data can lead to problematic and potentially damaging outputs.

Authors retain copyright and grant the journal right of first publication with the work simultaneously licensed under the *Creative Commons Attribution Non-Commercial 4.0 International License*.

Editorial Team

July 2026

Editor-in-Chief

Mahera Sarkar

 <https://orcid.org/0009-0009-7197-4375>

Managing Editors

Debarya Dutta

 <https://orcid.org/0009-0009-8632-413X>

Marine Ragnet

 <https://orcid.org/0009-0001-3343-3318>

Angy Watson

 <https://orcid.org/0009-0006-4488-4989>

Contents

Editorial	III.iii
Foreword <i>Marine Ragnet</i>	III.iv
Sustaining the Status Quo? Rethinking AI for Sustainability through Power and Infrastructure <i>Harry Collins</i>	131
Choosing the Path of Most Resistance: Metacognitive AI Literacy and Mitigating Epistemic Risk <i>Iman Khwaja</i>	142
Risk, Rights, and Regulatory Convergence: The Turkish Artificial Intelligence Law Proposal as a Model for Emerging-Economy AI Governance and the Prospects for a TRIPS-Style International Framework <i>Zeki Emre Kurt</i>	154
Resistance as a Technomoral Virtue for Our Times <i>Nitya Mandyam</i>	167
The Epistemic Frontier of Capitalism: A Theory of Epistemic Regression through Artificial Intelligence and Misinformation <i>James Rice</i>	177
Metacognitive Conversational AI: A Design Framework for Reflective Well-Being <i>Tianyi Yu</i>	188

Editorial

Welcome to the third issue of the *Cambridge Journal of Artificial Intelligence* (CJAI).

As the journal continues to grow, so too does the breadth and depth of the conversations it seeks to foster. What began as an ambition to create an interdisciplinary forum for the study of Artificial Intelligence has developed into a truly international community of scholars, students and practitioners committed to examining AI from diverse intellectual, cultural and professional perspectives. We are grateful to everyone who has contributed to this journey and are excited to present another issue that reflects the richness of this rapidly evolving field.

This edition marks a particularly exciting milestone for the journal through our collaboration with New York University's Peace Research and Education Program. Bringing together contributors from across institutions and disciplines, this partnership reflects a shared belief that understanding artificial intelligence requires engagement not only with technical innovation, but also with questions of governance, ethics, justice and human flourishing. We hope this collaboration will be the first of many that strengthen dialogue across academic communities and encourage new perspectives on the challenges and opportunities presented by AI.

The articles featured in this issue explore AI through the lens of power and responsibility. From metacognitive approaches to conversational AI and AI literacy to questions of technomoral resistance, governance, misinformation and sustainability, each contribution offers a thoughtful examination of how AI is reshaping the ways we think, make decisions and organise society. Together, these papers highlight that the most pressing questions surrounding AI are not solely about what these technologies can do, but about how

they ought to be designed, governed and integrated into our lives.

The continued success of the CJAI is only possible because of the generosity and dedication of so many individuals. I would like to extend my sincere thanks to our authors for entrusting us with their research and to our reviewers for their careful and constructive evaluations, which uphold the academic standards of the journal. I am equally grateful to our editorial team, whose professionalism, enthusiasm and commitment continue to shape the CJAI into a welcoming and rigorous platform for interdisciplinary scholarship. Finally, I would like to thank our colleagues at New York University for their collaboration in bringing this special issue to fruition.

As AI continues to evolve at an extraordinary pace, the need for thoughtful, evidence-based and interdisciplinary dialogue has never been greater. We hope the articles in this issue inspire further research, debate and collaboration, and we warmly invite readers to contribute to future editions of the journal as authors, reviewers or editors. Together, we can continue building a scholarly community that critically examines AI and its role in shaping our collective future.

With best wishes,



Mahera Sarkar
Founder & Editor-in-Chief

Foreword

This past June, in a university hall in Zomba, Malawi, a colleague held up a phone running a popular AI assistant and spoke to it in Chichewa. The room watched it stumble and give up. Nobody seemed surprised. The really telling part was the explanation that followed: these systems only learn from what they are fed, and everything left out is entirely invisible to them. For millions of Chichewa speakers, that invisibility is a daily reality. So are the indirect phrases a frightened woman in Thyolo might use to report a threat against her neighbour — the euphemisms people reach for when something is too dangerous or shameful to name directly. The workshop participants had spent the morning writing those very phrases on Post-It notes, simply because no system trained in Silicon Valley has ever encountered them.

I kept coming back to that room while reading the papers for this issue. It captures the exact question behind this call: where is power actually concentrating in AI, and what does that mean for the people living downstream of it?

When we first put out this call, a collaboration between the Cambridge Journal of Artificial Intelligence and NYU's Peace Research and Education Program, we expected substantial material analysis. We thought we'd get papers on water rights, data centres, supply chains, and labour. Instead, our contributors kept turning to something much more fundamental: power over what we know, how we come to know it, and whether we can still think for ourselves.

The six papers here group naturally into a few distinct arguments. The first is about epistemic control. One piece suggests that AI misinformation isn't a glitch, but a feature built into commercial ownership models. Another looks at AI literacy from a metacognitive angle, asking how we notice when a machine is doing our thinking for us, while a companion piece looks at the builder's side, offering a design framework that takes reflective well-being seriously rather than treating it as an afterthought.

Then there is the question of resistance, treated here not just as a public protest but as a quiet, deliberate skill. We have two papers arguing that the ability to say "no" to a system is worth teaching: one frames resistance as a "technomoral virtue," and the other argues for intentional friction as a form of literacy. The final two pieces look at governance from unexpected angles. One examines Turkey's proposed AI laws to see if emerging economies can dictate terms rather than just receive them, while the other shows how easily the language of "sustainability" can be co-opted to defend the very infrastructure it claims to critique.

That last problem, how the centres of power compromise themselves, came up in Zomba too, just on a much smaller canvas. During a workflow exercise on emergency reporting, the room completely stalled on one scenario: a woman reports a threat of violence, but the local chief who receives the alert happens to be the accused man's uncle. If a system can't route around a compromised local authority, it fails the people who need it most. It is the exact same structural challenge our governance papers wrestle with, just closer to the ground.

To be completely transparent about scale: this is a new journal, and this is a modest issue. We didn't get hundreds of submissions, and we aren't trying to pretend we did. What we have are six pieces of incredibly careful thinking, several from scholars and places that the field's usual citation loops completely ignore. But that was the whole point of partnering with a peace research program rather than sticking to traditional tech channels. The people most exposed to concentrated power are often the ones doing the sharpest thinking about it.

At the end of the day in Zomba, we took a photo of our "parking board", the flipchart where we left all the questions we couldn't resolve. We had told the room at the start that we didn't have the answers, just a shared desire to look for them. That feels like the right note for this issue. These papers don't wrap up the story of AI and power in a neat bow. Instead, they point carefully, from some unexpected directions, at

where we should be looking while that story is still being written.



Marine Collins Ragnet

*Head of Innovation Research, NYU Peace Research
and Education Program*

Adjunct Professor, New York University

*Managing Editor, Cambridge Journal of Artificial
Intelligence*