



Information Research – Vol. 31 No. 2 (2026)

A citizen's guide to artificial intelligence

Review of Zerilli, John. (2025). *A citizen's guide to artificial intelligence*.
Cambridge, MA: MIT Press. xxii, 209pp. ISBN 978-0-262-55392-6

DOI: <https://doi.org/10.47989/ir31265408>

A Citizen's Guide to Artificial Intelligence is a compact, wide-ranging attempt to give the lay reader a policy-aware account of AI. It does not treat AI as either magic or menace. Instead, it presents AI as a set of technologies embedded in institutions, that can improve public life, but only if they are understood, scrutinized, and governed properly.

The book's collaborative authorship is handled effectively. John Zerilli is the primary author, and contributors are listed. Zerilli's "voice", however, is carried through the book, since he wrote about half of it and incorporated the contributions in his own style. This gives the book a unified voice: clear and pedagogical. It feels written to guide readers through a complex field in a controlled progression, rather than to impress them with theory or technical detail. This is a real strength, because it allows the book to remain accessible while still covering substantial ground.

One of the book's most useful features is the way it is structured around thematic chapters, especially those dealing with bias, privacy, transparency, regulation, accountability, and the governmental use of AI systems. The chapter on bias is particularly important, because it moves the discussion away from abstract claims about "neutral" machines and toward the practical reality that AI systems inherit the values, assumptions, and distortions of the data and institutions that shape them. The book treats bias not as an incidental flaw but as a central design and governance issue. That is a valuable corrective to popular narratives that imagine algorithmic systems as objective simply because they are computational.

The discussion of regulation is similarly strong. Rather than treating regulation as an obstacle to innovation, the book frames it as a necessary condition for trustworthy AI. It explains why governance matters, what kinds of oversight are needed, and why the public should not rely on companies alone to police themselves. This section gives the book a distinctly civic character: it is not merely asking what AI can do, but what it should be allowed to do, under what rules, and for whose benefit. That emphasis makes the book especially relevant to readers interested in democratic institutions and public accountability.

Across the chapters, the book also does a good job of showing how AI systems raise questions of transparency, responsibility, and institutional power. It repeatedly returns to the idea that the central issue is not just technical performance, but the way AI is introduced into decision-making environments. That focus gives the book coherence. Even when the topics shift, the underlying concern remains the same: how can AI be deployed without eroding fairness, trust, and human oversight?

At the same time, the book was published in 2025, three years after the launch of ChatGPT, and yet it does not mention generative AI systems, large language models, or tools such as ChatGPT, even though these dominate much of the current public discussion. As a result, the book may feel incomplete to readers expecting an engagement with the most visible and controversial forms of contemporary AI.

The prose is accessible and the book appears designed for informed citizens rather than experts, and it succeeds in that goal by avoiding unnecessary technical detail while still respecting the reader's intelligence. The best chapters are those that combine conceptual clarity with policy relevance, showing how AI affects everyday institutions such as hiring, policing, welfare, finance, and public administration.

Overall, *A Citizen's Guide to Artificial Intelligence* is a thoughtful and well-organized and, even though it does not engage directly with generative AI, or LLMs, it remains valuable because it addresses the enduring questions that determine whether AI will serve public ends or merely private interests. It is a book that asks readers not just to understand AI, but to think politically about it.

Professor T.D. Wilson
April 2026

© [CC-BY-NC 4.0](#) The Author(s). For more information, see our [Open Access Policy](#).