Journal of Ethics and Emerging Technologies

IEET

Book Review

The Magician's Eye: A Review of David Eliot's *Artificially Intelligent: The Very Human Story of AI*

David Eliot, *Artificially Intelligent: The Very Human Story of AI.* (Toronto: University of Toronto Press, 2025)

Arvin Obnasca [1, 2,*]

[1] University of Sassari, Department of Humanities and Social Sciences, Italy; a.obnasca@uniss.phd.it
[2] BeEthical, France; arvin.obnasca@beethical.be
* Correspondence: arvin.obnasca@ethicists.ai

Abstract: David Eliot trained as a card magician before turning to AI research. His *Artificially Intelligent* sets out to democratise AI for general readers, and on its own terms it succeeds: the history is rich, the writing is clear, and the central argument that AI is socially constructed and democratically redirectable is held with conviction. This review reads the book alongside the regulatory architecture being built on the same theory of agency: the EU AI Act, the harmonised standards beneath it, the UNESCO and OECD ethics frameworks above, and the November 2025 Digital Omnibus provisionally agreed upon in May 2026. The argument here is additive, not corrective. Knowing the trick is not the same as breaking it. The real subject of magic is attention, not secrets. Alongside AI ethics and governance literature, the *magician's eye* points to what visibility cannot reach: the case for a code of practice from inside a profession that understands attention management.

Keywords: AI ethics, AI governance, AI standardisation, choice architecture, EU AI Act, GPAI Code of Practice, OECD Principles, UNESCO Recommendation

Citation: Obnasca, Arvin. 2026. The Magician's Eye: A Review of David Eliot's *Artificially Intelligent: The Very Human Story of AI. Journal of Ethics and Emerging Technologies* 36: 2. https://doi.org/10.55613/jeet.v36i2.239

Received: 17/05/2026
Accepted: 17/05/2026
Published: 01/07/2026

Publisher's Note: IEET stays neutral with regard to jurisdictional claims in published maps and institutional affiliations.

David Eliot's *Artificially Intelligent* begins with a convincing admission. Before he became a researcher in artificial intelligence (AI), he was a magician who specialised in cards. Following this introduction, there's a persistent metaphor at play: AI has power over us because we don't know how it works, and when you pull back the curtain, you see the trick is "actually mundane" (Eliot, 2025, Chapter 1).

It's a likeable opening, no doubt, more interesting than it first appears.

For two centuries working magicians have known that explaining does not end the effect. And that's what Penn and Teller have built their career on. They walk you through the trick and do it anyway, and you still don't see. What attention does is stage magic, when it occurs.

The book is more pointed on magic than its opening suggests. In the first chapter, the author names the architecture: the audience sees fifteen percent of the show; the real work happens in the hidden eighty-five percent. Magicians spend their lives perfecting moves that no one is supposed to notice. Then Eliot turns. AI, according to his reading, is the opposite of magic: the magic show keeps its secrets; AI's (secrets) should be revealed. The professional understanding of how magic affects its audience goes out the window when the discussion turns to AI.

What does a magician know? Not, mainly, secrets. Apollo Robins is famous for stealing a watch from a man who has just been told he is about to take it (Robbins, 2013). The technique is framing, suggestion, and where you point the eye while the hand is busying itself with another task. Macknik, Martinez-Conde and Blakeslee (2010) present magic as applied research in steering attention. The magician controls what you perceive.

This is a hilariously good description of how recommendation engines, ad targeting, platform design, and nudges actually work. None really covers anything up. Data practices are disclosed; retention windows live in the policy; training methods appear in research papers. Nissenbaum (2011) identified the dynamics over a decade ago: a transparency paradox where more disclosure leads to less understanding. Stage magic supplies the active half. The systems are not merely too complicated to decipher; they are engineered to divert the eye, shape the choice, and create a sense of free decision when the alternatives have already gone elsewhere. Similar formulations arise from the literature on dark patterns and on choice architecture. The work on managing attention, not hiding.

Eliot's book goes another way through the same material. This review goes along with it.

Read the book with a magician's question, not "is the analysis correct," but "what is the performance doing." Someone who reads *Artificially Intelligent*, worried about being controlled by AI, puts it down, feeling empowered to control the story. The book effects that change, but how?

Most of the shift is due to structural choices. In fact, the book is overtly theatrical: Eliot concludes the prologue by saying that he will tell the story of the rise of AI "as it is best understood: as a story" (Eliot, 2025, Prologue), and the chapters are structured in three Acts. He names the structural forces: Google, Facebook, and the surveillance state. The last chapter is a call to action. None of those moves gives the reader a real hand in shaping AI; they give the feeling of having a hand in it. The feeling is the rabbit out of the hat, as in magic.

This is a familiar framing. What's interesting here is the overt staging. Eliot does not conceal that he is performing. What is missing is the transition from staging-as-vehicle to staging-as-object. The book uses that structure of the theatre to tell its story; it does not turn that structure in on itself.

Several passages come close to the magician's subject without quite landing on it. In the chapter on the new computer, Eliot describes Weiser's vision of ubiquitous computing in reverse: machines meant to serve human needs may instead be the nightmare he was trying to prevent. In the Vision Pro chapter, people are "cyborgs of sorts" with ambient sensors that record eye movement, brain activity, and heartbeats (Eliot, 2025, Chapter 18). In the chapter "The Ubiquitous Machine," the YouTube algorithm is described as not just deciding what viewers watch, but also training creators (Eliot, 2025, Chapter 21).

These are the right occasions for a magician's reading. It's not deeply opaque. The Vision Pro sensors are not secret. The YouTube creator economy has been well documented. That's what you're seeing. Misdirection on a scale. The system points the eye one way, and shapes the person another. The cyborg of sorts is not fooled about some hidden mechanism; he is briefed and still doing what the system wants him to. That's what magic looks like on someone who read the manual.

Eliot does not dwell upon such moments. The first chapter establishes a transparency framing that the reader can follow forward.

Let us study it further and see where this faith comes from. The ethics frameworks were already in place. Transparency and explainability are named as fundamental by UNESCO's Recommendation, in addition to human oversight, fairness, and dignity (UNESCO, 2021); OECD's AI principles do similar work for forty-seven adhering governments (OECD, 2024). Both echo Eliot: tell people what the system is, give the literacy to understand it, and trust will come. They also have the vocabulary beyond disclosure, such as "human autonomy," "human dignity," and "human-centred values." Mittelstadt (2019) argued that principle convergence does not result in ethical AI because the principles do not hold up under implementation. Floridi (2019) calls this as a dynamic of 'ethics shopping,' where principles are attached onto behaviours after the fact, rather than the other way round. By the time of audit, most of the moving has been done by transparency; dignity and autonomy remain in the preamble.

So this is where a magician's reading does the most useful work. The European Union Artificial Intelligence Act (EU AI Act), in force since August 2024, inherits the vocabulary as well as the thinning. It can be assessed to be organised around two competing paradigms that a magician would instantly recognise as different theories of how harm occurs. Most readers see an Act. But there are two ideas in it.

Paradigm 1: Disclosure. Looking at Article 4 (AI Literacy) and Article 50 (Transparency Obligations) vis-à-vis Eliot's idea gives us the notion of "knowing is half the protection."

Paradigm 2: Prohibition. Examining Article 5 (Prohibited AI Practices) and its lomg list presents us with the realities of limits that "just knowing is not enough."

The second paradigm has the right instinct and the wrong calibration, a magician would say. Relating to stage magic, which works around "managed attention regardless of disclosure" and not simply "hidden technique," it has the right instinct. But insofar as the Act's threshold is a "subliminal technique beyond a person's consciousness" that "materially" distorts behaviour and causes "significant harm" (Article 5(1)(a)), it has the wrong calibration.

In a complementary argument on Article 50, we might even ask: To what extent is "disclosure" useful (or useless) if it does not protect the users from what happens next?

There is another tension that surfaces on Article 14 (Human Oversight), where high-risk AI systems have to leave a room for human intervention. A magician would look at it as a performance with a known ending. How a system presents an output affects whether the viewer engages. Confidence scores come across as authoritative. Default accept buttons are in the natural reading path; investigate options are behind a second click. Add time pressure, and the documented response is automation bias. Article 14 calibrates against an absence of an override, not an unused one.

Another layer is behind the Act. Compliance is delegated to harmonised European standards that are being developed by CEN-CENELEC's Joint Technical Committee 21 (Regulation (EU) 2024/1689, Article 40; CEN-CENELEC, 2025). Let us see which came first. The *Quality Management System* for AI Act regulatory purpose (prEN 18286), built around Article 17 of the AI Act, was the first harmonised AI standard to reach formal public enquiry in October 2025. Published ahead of it internationally is the *AI Management System* (ISO/IEC 42001:2023) for AI governance. Other substantive standards are still being written. Even then, we notice that practically, compliance means documenting the processes.

To the extent that standards translate the vocabulary of the Act into something measurable, we can say that they do real work where measurement applies, and less where it doesn't.

Running parallel to the technical requirements discussion is the timeline development. Significant adjustments have since taken place between the November 2025 Digital Omnibus on AI proposal of the Commission (European Commission, 2025) and the 7 May 2026 trialogue agreement reached between the Council and Parliament negotiators (Council of Europe Union & European Parliament, 2026). The Omnibus revises deadlines, documentation, literacy obligations, and the threshold for processing sensitive data. It does not return to disclosure-as-protection, the tuning of Article 5 for overt attention management, or whether an audit can see what an audit was not designed to see. The EU took the Act apart and put it back together again without any mention of the issue of managed attention.

Peel back these layers of the book, and the deeper claim of this review is revealed. *Artificially Intelligent*, the ethics frameworks above it, the Act's disclosure paradigm, and the standards below are doing the same work. Each of them gives a sense of being able to do things. None provides it, at least not yet. The people Eliot describes (the YouTube creator remaking their work for the algorithm, the Vision Pro user whose biometrics are read back, the cyborg of sorts at the end of the Apple chapter) are harmed in ways that disclosure and audit instruments above do not fix. Their problem is downstream of disclosure, in what the system has done with their attention by that point.

And that's where conversation ethics could come in handy. More than disclosure, UNESCO and the OECD have vocabulary. There is a code of practice as well. The General-Purpose AI Code of Practice was drafted by independent experts with input from stakeholders (European Commission, 2025a) and published in July 2025. It is about transparency, copyright, and being safe for general-purpose model providers. What it does not do is get to the deployment surface where the system interacts with the user. That's the level a magician would be interested in. These are things a serious performer will not do to a paying audience, even when there is consent. What's interesting is that it's a code in that spirit, from people who understand attention management, applied to systems whose business models match misdirection at scale. Eliot is in a good position to begin that dialogue.

None of this is a complaint. With a voice warm and authoritative, an impressive sweep of history, and a rare political conviction, *Artificially Intelligent* is indeed a strong general-reader introduction. It is the right book for a curious friend.

The *Magician's Eye*, then, implies a companion book. One that asks a different question: not what AI hides that we should uncover, but what AI asks us to see as we look. *Artificially Intelligent* waves in that storyline. Eliot may be the man to take it.

Funding: This book review received no external funding.

Conflicts of Interest: The author declares no conflict of interest.

Acknowledgement: All views and interpretations expressed in this review are the author's own. Grammarly Premium supported textual refinement and readability improvements.

References

1. CEN-CENELEC. (2025, October 23). *Update on CEN and CENELEC's decision to accelerate the development of standards for artificial intelligence*. https://www.cencenelec.eu/news-events/news/2025/brief-news/2025-10-23-ai-standardization/
2. Council of the European Union & European Parliament. (2026, May 7). *Artificial intelligence: Council and Parliament agree to simplify and streamline rules* [Press release]. https://www.consilium.europa.eu/en/press/press-releases/2026/05/07/artificial-intelligence-council-and-parliament-agree-to-simplify-and-streamline-rules/
3. Eliot, D. (2025). *Artificially intelligent: The very human story of AI*. Aevo UTP / University of Toronto Press.
4. European Commission. (2025a, July 10). *General-Purpose AI Code of Practice*. https://digital-strategy.ec.europa.eu/en/policies/contents-code-gpai
5. European Commission. (2025b, November 19). *Proposal for a regulation amending Regulations (EU) 2024/1689 and (EU) 2018/1139 as regards the simplification of the implementation of harmonised rules on artificial intelligence (Digital Omnibus on AI)* (COM(2025) 836 final). https://digital-strategy.ec.europa.eu/en/library/digital-omnibus-ai-regulation-proposal
6. European Parliament & Council of the European Union. (2024). Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Official Journal of the European Union*. https://eur-lex.europa.eu/eli/reg/2024/1689/oj
7. Floridi, L. (2019). Translating principles into practices of digital ethics: Five risks of being unethical. *Philosophy & Technology*, 32(2), 185-193. https://doi.org/10.1007/s13347-019-00354-x
8. International Organization for Standardization & International Electrotechnical Commission. (2023). *ISO/IEC 42001:2023 Information technology, Artificial intelligence, Management system*. ISO. https://www.iso.org/standard/42001
9. Macknik, S. L., Martinez-Conde, S., & Blakeslee, S. (2010). *Sleights of mind: What the neuroscience of magic reveals about our everyday deceptions*. Henry Holt and Company.
10. Mittelstadt, B. (2019). Principles alone cannot guarantee ethical AI. *Nature Machine Intelligence*, 1, 501-507. https://doi.org/10.1038/s42256-019-0114-4
11. Nissenbaum, H. (2011). A contextual approach to privacy online. *Dædalus*, 140(4), 32-48. https://doi.org/10.1162/DAED_a_00113
12. Organisation for Economic Co-operation and Development. (2024). *OECD AI Principles* (updated May 2024). https://oecd.ai/en/ai-principles
13. Robbins, A. (2013). *The art of misdirection* [Video]. TED Conferences. https://www.ted.com/talks/apollo_robbins_the_art_of_misdirection
14. United Nations Educational, Scientific and Cultural Organization. (2021). *Recommendation on the Ethics of Artificial Intelligence*. UNESCO. https://www.unesco.org/en/articles/recommendation-ethics-artificial-intelligence