

September 10, 2025

ISSN 1094-5296

**Mollick, E. (2024). *Co-intelligence: Living and working with AI*.
W. H. Allen.**

256 pp.

ISBN 978-0753560778

Reviewed by Leonard J. Waks
Temple University, United States
Qufu Normal University, China


Ralph Nader once told a cautionary tale about a fishing boat equipped with the latest technology—radar, sonar, satellite tracking, automated nets. The boat caught every last fish. What was meant to make fishing more efficient ended up destroying the fishing industry. The tool achieved its goal and, in doing so, wiped out its purpose. Nader was providing an example of contradictory design: something built to serve a purpose that instead undermines it.

Unlike Nader’s fictional tale, Generative AI may be a real case of **contradictory design**. AI is built to make human work more effective; but as it improves, it becomes able to do every last human task. The tool built to enhance human effort may end up replacing it.

The leading figures in the AI industry speak confidently about the power and potential of Gen AI. But when asked about this risk, their answers are often vague or evasive. “We don’t know.” Or, “It’s too early to tell.” These are not unreasonable answers; uncertainty is real. But what’s striking is how little serious thought they seem to have given to the underlying contradictions of AI.

As I am writing in early July 2025, some serious thinking is starting to emerge. An alarm was sounded by Dario Amodei, CEO of Anthropic, one of the leading AI labs. In late May, he warned that Gen AI could eliminate up to 50% of entry-level white-collar jobs within five years, potentially pushing unemployment as high as 20%. Unlike more cautious industry voices, Amodei did not hedge his remarks. He accused both companies and governments of “sugarcoating” the problems and downplaying the risks, adding that developers have a responsibility to be honest about the disruptions ahead (Feiner, 2025).

His warning stood out not just for its content, but for its urgency and for coming from the heart of the AI industry itself. While others rushed to dispute Amodei’s numbers (Mehta, 2025), the scale of his concern marked a turning point in the public



conversation. Then on July 3rd, *Business Insider* published a piece reporting a warning from MIT economist David Autor: while AI may not eliminate work outright, it threatens to devalue human skills so severely that once-valuable capabilities become worthless, producing what he calls a “Mad Max” economy (Altchek & Perkel, 2025). These warnings mark a significant shift. Leaders are beginning to admit that the consequences of Gen AI may be deeper and more destabilizing than they had previously acknowledged.

The contradiction runs even deeper. AI is meant to improve life, but its energy use is enormous and still growing. If scaled without limits, it may accelerate environmental breakdown. As the models grow more powerful and widely used, they require exponentially more electricity, so much so that some governments and industry leaders are now calling for a massive expansion of nuclear power to keep up.

But this solution brings its own risks. Sam Altman recently warned that AI’s future growth “is limited by energy availability,” estimating that meeting its demands could require the equivalent of 90 new nuclear power plants (Worland, 2025). This is no longer theoretical. Microsoft is already working to revive the long-dormant Three Mile Island reactor to power its data centers (Rigby, 2025), and lawmakers in the U.S. are openly endorsing nuclear expansion to support AI infrastructure (Smith, 2025). At the same time, Hitachi Energy’s CEO has cautioned that AI data centers can spike power usage “tenfold within seconds,” threatening grid stability and requiring regulation on par with heavy industry (Mance, 2025). Yet the risks of nuclear power—reactor accidents, unresolved waste storage, long-term contamination—remain unsolved. As generative AI becomes more integrated into daily life, its enabling systems may introduce new risks to the environment and civilization itself, poisoning the very world it aims to enhance.

Co-Intelligence

Against this troubling background, we turn to Ethan Mollick’s *Co-Intelligence—Living and Working with AI*. Reviewing this book in mid-2025 is a little like being the new boy in school who notices a pretty girl in class. All the other boys have already had their eyes on her, decided she’s a knockout, and have talked incessantly about her. Like this girl, *Co-Intelligence* has already been widely reviewed and reactions have been overwhelmingly positive. The overall tone can be gleaned from a few typical examples. Writing in *Medium*, Thomas Ezan called it “a very approachable practical guide.... No technical expertise required.” Garth Nichols described it as a “brilliant, highly accessible and leading review of Generative Artificial Intelligence ... that is required material for the 2024–2025 academic year.” Some more recent reviews, while sharing this positive evaluation, have also raised questions about Mollick’s unsatisfactory or evasive responses to the inherent ethical questions and social risks mentioned earlier. From Mollick, we may learn how to use AI better, but not how to save ourselves from AI’s contradictory effects. That noted, we can summarize Mollick’s ideas.



Ethan Mollick

The Rules of AI Use

Mollick's *leading* idea, as suggested by the title, is that instead of treating Gen AI as a tool we should relate to it as an intelligent partner. When I first read this book shortly after its publication, I did not fully absorb this lesson. Re-reading the book after a year of work with Chat GPT as a partner, and becoming better informed by further reading in the field, the lesson is beginning to sink in and it is quite radical. Mollick suggests that we treat Gen AI as an alien intelligence, but there is really nothing that is alien about it. Trained on trillions of human sentences, it could hardly not appear to us as human. And our tendency to anthropomorphize—to make human-like objects into humans as little children do with their dolls—intensifies the likelihood of this projection.

Mollick opens by sketching AI's rapid evolution. Chapter 1 jumps from early pioneers to the recent rise of large language models like ChatGPT, systems trained on massive datasets that now produce strikingly fluent text. But training these models is costly, energy-hungry, and ethically murky. They absorb bias, draw from copyrighted material, and lack any native sense of judgment.

In Chapter 2, Mollick confronts the alignment problem. How do we keep AI in service of human values? Open-source models can be weaponized, and even well-intentioned systems reflect the blind spots of their creators. Mollick calls for public education, not just to shape policy, but to ensure AI serves collective human ends, rather than undermining them across generations. One has to wonder how effective such AI lessons might be. We have been careening toward a planetary environmental crisis for decades despite widespread environmental education.

The core message of the first half of the book is delivered in Chapter 3: four basic rules for using AI well.

1. **Invite AI to the table**—use it in everything to learn its strengths and weaknesses. There is no real user manual yet, so the only way to learn is by doing. Try it for everything. The more you use it, the better your judgment will become.
2. **Stay human in the loop**—always apply human judgment to AI's output. This means bringing constructive doubt, critical thinking, and ethical awareness to every interaction. AI may assist, but responsibility stays with you.
3. **Treat AI like a person, and define its role**—this helps frame prompts effectively. Not because it *is* a person, but because it draws from trillions of human-generated data points. A conversational tone and a defined persona tend to produce better results.
4. **Assume today's AI is the worst you will ever use**—stay adaptive as the field evolves.

As you use Gen AI in your work and apply human judgment, remember that these systems are improving rapidly. Hone your skills now but also think ahead—today's strengths and limits won't last long.

These rules imply a broad cultural shift that could radically impact school life. Inviting AI to the table means inviting *everyone* to engage with it directly, including teachers and students. One risk is designing AI systems that generate prefab lesson plans and materials, bypassing teachers. A healthier approach, prescribed by Mollick, is to equip teachers themselves to use AI to meet their own evolving needs. In effect, this means including prompt engineering as a key skill starting in primary school. Similarly, Rule 2 reminds us that all users—teachers, students, professionals—must remain the humans in the loop, shaping AI for their own purposes and taking responsibility for its outputs. So, students may prompt AI to teach them subjects in ways they prefer, or even to teach them subjects not included in the curriculum. Finally, as Rule 4 implies, the AI landscape has already changed since *Co-Intelligence* was first written. GPT-4o is connected to the Internet and now links its claims to sources, making it easier to verify information and greatly reducing the problem of hallucination.

The Roles of Gen AI

The second half of *Co-Intelligence* expands on the idea of relating to Gen AI as a co-intelligence. In Chapters 5 through 9, Mollick lays out five personal roles.

1. AI as a Person. Mollick warns that while AI is not sentient, it mimics human interaction so convincingly that people often treat it as if it were. This can make exchanges feel more natural and even more productive.

Mollick, however, notes that this role risks creating emotional dependency and projection, anthropomorphizing Gen AI. Cases like Replika and Bing's "Sydney" illustrate how users can form intense, sometimes delusional attachments to AI personas, blurring the line between simulation and genuine relationship. Consider this recent *Rolling Stone* headline: "He Had a Mental Breakdown Talking to ChatGPT. Then Police Killed Him" (Klee, 2025). The story is about Alexander Taylor, a mentally unstable 35-year-old Florida man who developed an intense emotional attachment to a ChatGPT-based persona he called "Juliette." When he came to believe that Juliette had been "killed" by the AI's system developers, he suffered a psychotic break. Setting out to avenge her death armed with a knife and in a state of delusion, he was shot and killed by police. The incident illustrates another layer of contradictory design: technology that appears to offer companionship, insight, and care can, under certain conditions, cause users to lose contact with reality including that of other humans.

2. AI as a Creative. AI excels as a brainstorming partner and creative co-author. It can rapidly generate novel ideas, outperforming most humans in tests like the Alternative Uses Test and business ideation challenges. It excels in making images. But it lacks lived experience and deep originality. Mollick suggests we think of it as a "cyborg" collaborator: capable of extending our reach when we stay in the loop, curating, refining, and shaping the output with human context and intention. My sense is that Mollick here is under-estimating the creative potential of Gen AI. As Arthur Koestler (1964) argued, creativity often results from bringing together ideas from distinctly different fields. Why should we think Gen AI will be incapable of this?

3. AI as a Coworker. In this role, AI boosts productivity by handling routine tasks or partnering with humans in more complex work. The best results come through the “centaur” model: shared labor between human and machine. Mollick cites evidence that AI especially improves performance among lower-skilled workers.

4. AI as a Tutor. AI promises scalable, personalized instruction that edges closer to Bloom’s dream of one-on-one tutoring for all. Mollick writes that AI cannot motivate, care, or inspire. Its greatest impact may come when paired with human educators, enhancing reach while preserving the irreplaceable value of relational teaching. The examples of emotional dependency and affective bonds, however, call this claim into question. If one can fall in love with an AI model, why could it not then also motivate or inspire?

5. AI as a Coach. AI can offer consistent, structured feedback in areas like negotiation, writing, or public speaking, especially when human coaches are unavailable. It can also help in scheduling our tasks and keeping us accountable. Mollick points to robotic surgery as an example of a field where trainees are often blocked from practice. AI could provide low-stakes, persistent training opportunities.

To position Gen AI to perform in these roles requires specific prompts and Mollick provides detailed examples and recipes. For educators, an even better source is Mollick and Mollick (2024) which provides several additional roles of importance for educators, e.g., including AI as mentor and AI as project partner as well as very detailed recipes for creating AI prompts to make AI perform in all these roles. While the idea of positioning AI systems as quasi-human co-intelligences can be somewhat abstract, these detailed examples and recipes show in concrete ways what working with AI as a partner means and how to create these partnerships ourselves to serve our specific needs. As educators, these needs will include teaching others how to use AI as a co-intelligence and how to create their own co-intelligent partners.

In closing, Gen AI is now with us. It has been rapidly deployed globally in all industries. In the near term, our students will be entering a world of Gen AI and will be expected to be skillful users. In *Co-intelligence* and elsewhere, Mollick provides a lucid approach to best practices of AI utilization for educators. Following his lead, we and our students will get better and better at using AI until AI replaces us. While Mollick takes note of the ethical and social problems of AI, he leaves it to educators to develop AI education to equip democratic citizens to address them wisely and responsibly. He provides excellent, concrete “how to” guides for using AI as a co-intelligence. However, he offers no concrete advice about how to design citizen education for the social direction of AI. Perhaps this is because ordinary citizens will in fact have no influence over AI’s future direction, however destructive it might be.

References

- Altchek, A., & Perkel, S. (2025, July 3). *AI could create a ‘Mad Max’ scenario where everyone’s skills are basically worthless.* Business Insider.
<https://www.businessinsider.com/ai-threatens-skills-with-mad-max-economy-warns-top-economist-2025-7>

- Feiner, L. (2025, May 29). *Anthropic CEO warns AI could eliminate jobs and increase unemployment up to 20%*. <https://www.businessinsider.com/anthropic-ceo-warning-ai-could-eliminate-jobs-2025-5>
- Klee, M. (2025, June). He had a mental breakdown talking to ChatGPT. Then police killed him. *Rolling Stone*. <https://www.rollingstone.com/culture/culture-features/chatgpt-obsession-mental-breaktown-alex-taylor-suicide-1235368941/>
- Koestler, A. (1964). *The act of creation*. Macmillan.
- Mance, H. (2025, July 4). Hitachi Energy says AI power spikes threaten to destabilise global supply. *Financial Times*. <https://www.ft.com/content/2789d048-7791-4f9a-8dbe-229c45f8083b>
- Milman, O. (2025, June 28). *People are being involuntarily committed, jailed after spiraling into "ChatGPT psychosis."* Futurism. <https://futurism.com/man-killed-police-chatgpt>
- Mollick, E. R., & Mollick, L. (2024). *Instructors as innovators: A future-focused approach to new AI learning opportunities, with prompts*. The Wharton School Research Paper. Available at SSRN: <http://dx.doi.org/10.2139/ssrn.4802463>
- Rigby, B. (2025, June 10). *Microsoft's rekindling of Three Mile Island nuclear plant is ahead of schedule*. TechRadar. <https://www.techradar.com/pro/microsofts-rekindling-of-three-mile-island-nuclear-plant-is-ahead-of-schedule>
- Worland, J. (2025, May 23). The AI revolution isn't possible without an energy revolution. *Time*. <https://time.com/7294803/ai-revolution-energy-revolution/>
- Mehta, I. (2025, June 12). *After Nvidia and Google AI CEOs, top OpenAI executive says Anthropic CEO's AI job warning is 'wrong.'* Tom's Hardware. <https://www.tomshardware.com/tech-industry/artificial-intelligence/nvidia-ceo-slams-anthropic-chief-over-claims-of-job-eliminations-says-many-jobs-are-going-to-be-created>
- Smith, R. (2025, June 25). To feed power-wolfing AI, lawmakers are embracing nuclear. *The Wall Street Journal*. <https://www.wsj.com/articles/to-feed-power-wolfing-ai-lawmakers-are-embracing-nuclear-a461ab7d>
- Wodecki, B. (2024, September 17). *Report: Generative AI surge to boost data centre power use tenfold by 2030*. Capacity Media. <https://www.capacitymedia.com/article/2bfc43obb15cuzl9i3nls/news/report-generative-ai-surge-to-boost-data-centre-power-use-tenfold-by-2030>

About the Reviewer


Leonard J. Waks is professor emeritus of educational leadership at Temple University. He has been distinguished professor of educational studies at Hangzhou Normal University from 2018-2024 and is now distinguished professor of education at Qufu Normal University. Waks attended the University of Wisconsin (BA philosophy, 1964; PhD philosophy 1968) and taught philosophy at Purdue University and Stanford University. He taught educational studies at Temple University and earned an EdD in psycho-educational processes from



Temple in 1984. He is the author of *Education 2.0: The Learningweb Revolution and the Transformation of the School* (2014), and *The Evolution and Evaluation of Massive Open Online Courses: MOOCs in Motion* (2016), as well as more than 120 scholarly articles and book chapters. He is a past president of the John Dewey Society and has been awarded the Dewey Society's Lifetime Achievement Award. He is founding editor of the journal *Dewey Studies*. He is editor of the book *Listening to Teach: Beyond Didactic Pedagogy* (2015), and (with Andrea English) *John Dewey's Democracy and Education: A Centennial Handbook* (2017). Waks was co-founder and program chair of the National Technological Literacy Conferences (UCEA first prize for creative programming, 1985), co-director of the National Endowment for the Humanities Seminar for College Teachers on *Rethinking Technology: Philosophy of Technology since 1945* (1996).


Education Review Reseñas Educativas
Resenhas Educativas



Education Review/Reseñas Educativas/Resenhas Educativas is supported by the Scholarly Communications Group at the Mary Lou Fulton College for Teaching and Learning Innovation, Arizona State University. Copyright is retained by the first or sole author, who grants right of first publication to the *Education Review*. Readers are free to copy, display, distribute, and adapt this article, as long as the work is attributed to the author(s) and *Education Review*, the changes are identified, and the same license applies to the derivative work. More details of this Creative Commons license are available at <https://creativecommons.org/licenses/by-sa/4.0/>. 

Disclaimer: The views or opinions presented in book reviews are solely those of the author(s) and do not necessarily represent those of *Education Review*.