

Book review: Ethan Mollick. *Co-Intelligence: Living and Working with AI*. Portfolio/Penguin, 2024.

Michael Ridley 
University of Guelph

Books about where AI is taking us are being released almost as rapidly as new LLM models. While Yogi Berra famously said, “it’s tough to make predictions, especially about the future,” this is exactly what generative AI does. With *Co-Intelligence: Living and Working with AI*, Ethan Mollick steps up to the plate.

Mollick is a faculty member at Wharton specializing in entrepreneurship and innovation. He is well known for his engaging and valuable newsletter *One Useful Thing*. The central premise of the book is important: with advanced AI we will have not just an intelligent partner or a cognitive prosthetic but an autonomous “co-intelligence” that “potentially replaces human thinking.” Notice, however, the caveat. While remaining optimistic, a considerable part of the book wrestles with the implications of that.

This book is intended for a popular audience. It is readable, accessible to the non-expert, often humorous, filled with instructive anecdotes, and supported albeit lightly with reference to the underlying research. For those working in and around this field, it will not advance your understanding or provide new insights. However, is this a book you could recommend to someone who wants to understand what generative AI is, what it does, and what it could do? Yes, but act quickly. Like so much about AI, it’s best before date is looming. The valuable parts of this book attempt to rise above the fleeting moment.

Mollick devotes two chapters to AI as “alien minds.” First, he discusses creating them and he raises the challenge of aligning with them. He usefully expands on these by looking at a series of roles AI will take on: person, creator, coworker, tutor, and coach. However, at the outset, in identifying the co-intelligence of AI as an unprecedented “new thing,” Mollick says that until now “we have never built a generally applicable technology that can boost our intelligence.” This ignores, tellingly, the alphabet and writing, and gives us a whiff of the technology cheerleading that surrounds this book.

Mollick tries to be positive about AI while still acknowledging the serious challenges and even dangers (although how many times do we need Bostrom’s paper clip story as a

cautionary tale?). How we will navigate and mitigate those challenges is a weakness. Mollick says regulations are not going to be enough and instead “the path forward requires a broad societal response” with “norms and standards” and “informed citizens” with everyone working together. Rather than focusing on AI catastrophes he unearthed J. R. R. Tolkien’s neologism “eucatastrophe” which describes a happy ending, a good catastrophe. Ultimately Mollick says, “we are going to decide” about what AI does and the implications. Of course, who the “we” are is significant and largely undiscussed.

The most valuable part of this book is Mollick’s “four rules for co-intelligence:”

1. Always invite AI to the table
2. Be the human in the loop
3. Treat AI like a person (but tell it what kind of person it is)
4. Assume this is the worst AI you will ever use

Whatever the deficiencies of his speculations, these are (mostly) wise observations and guidelines. Encouraging us to purposefully engage with AI and to bring critical reflection to the table are important recommendations. Assuming that the AI you are using is the worst one you will ever use recognizes the pace of change and tempers the rush to judgment. It also is an acknowledgement that we can make things better, implicitly referring us back to rules #1 and #2. However, treating AI like a person misses the mark in a very serious way. Despite his prevarications about anthropomorphizing, machines are nothing like people and do not deserve even the metaphorical status he grants them.

Arising from his observations, Mollick outlines four scenarios for the future. AI now could be as good as it gets (unlikely); there could be slow growth in performance (not the experience so far); exponential growth could happen (continuing what we have been experiencing so far); or, ominously, “the machine God” (the rise of AGI and superintelligence). The latter is, depending on where you stand and to continue Mollick’s metaphor, either the apotheosis or the apocalypse. I find all this very shallow and unhelpful. As a bit of an aside, for those in the LIS field, I would instead recommend the *2035 Scenarios: AI-Influenced Futures in the Research Environment* recently issued by the Association of Research Libraries (ARL) and the Coalition for Networked Information (CNI).