BOOK REVIEW

Thinking About Co-Intelligence

Co-Intelligence: Living and Working with AI
Ethan Mollick (Portfolio | Penguin Random House 2024),
234 pages

Katrina Robinson, rev'r*

Generative AI can feel like a runaway train. Even if you were one of the lucky ones who saw the train coming and managed to clamber aboard, the breakneck pace and uncharted destination make it difficult for you to keep your foothold. For those who watched the train race by in shock, looked the other way in denial, or were caught blissfully unaware, the train is now so far in the distance that catching up to it can seem impossible. This reality is problematic for members of the bench, bar, and academy, as AI is poised to have profound effects on legal practice and education. Fortunately, a new book, *Co-Intelligence: Living and Working with AI*, offers guidance for both groups—the experienced riders and the would-be passengers.

The book's author, Ethan Mollick, is a professor of innovation and entrepreneurship at the Wharton School of the University of Pennsylvania. He describes his research as studying "how to teach people to become more effective leaders and innovators" and "how technologies are *used*."[1] Working in these areas led him to be an early enthusiast for AI's applications in education and business. In November 2022, he began writing a free newsletter, *One Useful Thing*, to provide a "research-based view on the implications of AI."[2] His newsletter now has a significant following, and his recent book is likely to enjoy similar success.

* Assistant Professor of Law, Cornell Law School.

1 Ethan Mollick, LinkedIn Profile, https://www.linkedin.com/in/emollick/ (last visited May 15, 2024); Ethan Mollick, Co-Intelligence: Living and Working with AI xix (2024).

2 Ethan Mollick, *Welcome to One Useful Thing*, One Useful Thing, https://www.oneusefulthing.org/about (last visited May 16, 2024); Ethan Mollick, *How to . . . Be More Creative*, One Useful Thing (Nov. 10, 2022), https://www.oneusefulthing.org/p/how-to-be-more-creative.

Drawing on his experience teaching undergraduate and MBA students, his experiments with previous iterations of AI tools, and his active involvement in emerging research on practical uses for AI, Mollick has crafted a thought-provoking and accessible book about AI. The book proceeds in two Parts: Mollick describes Part I as answering the basic question of "What is AI?" so that readers have a basis for thinking about how to work with AI systems, and Part II as discussing "how AI can change our lives by acting as a coworker, a teacher, an expert, and even a companion."[3]

Three Sleepless Nights

The book opens with a cautionary note: Getting to know AI will cost the reader at least three sleepless nights. For Mollick, the insomnia began shortly after the release of ChatGPT in November 2022. He had typed a paragraph-long prompt asking the bot to fill the role of a teacher in creating a detailed negotiation simulation, providing feedback on his performance in the simulation, and assigning him a grade. Simulations like the one he described in the prompt are a key feature of Mollick's own pedagogy and research. In fact, as Mollick shares, for the last five years, he and a team of collaborators have been developing "elaborate digital experiences" to simulate the business world and teach relevant skills like negotiation.[4] But according to Mollick, in a matter of minutes, ChatGPT "did 80 percent of what took our team months to do."[5] The bot's response to Mollick's prompt was imperfect, but quite good.

After establishing himself as someone who also stands to gain and lose something with AI's advances, Mollick invites readers "on a tour of AI as a new thing in the world, a co-intelligence, with all the ambiguity that the term implies."[6]

A Nebulous Term

Chapter 1 explains that "AI" is a nebulous term that has meant different things to different people at different times. For readers who might have a narrow definition in mind, this opening note is clarifying. This framing may also be strategic as it gives Mollick some leeway in

3 Mollick, *supra* note 1, at xx.

4 *Id.* at xiv.

5 *Id.* at xiv–xv.

6 *Id.* at xix.

deciding how to tell the story of AI's development. The first plot point on his timeline turns out *not* to be 1956, when John McCarthy of MIT coined the term, but rather, 1770, when the first mechanical chess computer was invented and began touring the world. This is a surprising starting point because, as Mollick reveals, the machine was eventually exposed as a fiction—a human chess master hid inside the gears, controlling its moves in every game! But by including this vignette on his timeline, Mollick hopefully assures the reader that he intends to offer a balanced view of AI in the pages to come. In fact, throughout the book, Mollick pauses to acknowledge relevant ethical lapses and other problematic moments in AI's development.

Mollick doesn't spend too much time discussing old technologies, though. Instead, he provides a helpful gloss on the "boom-and-bust cycles" of AI development, explaining how, like other technologies, AI research and development rises and falls with the excitement of investors.[7] Along the way, he introduces key terms and concepts that readers might have heard in discussions about AI like "artificial neural networks," "machine learning," "supervised learning," and "algorithmic decision-making."[8] Techy readers will fly through these ten pages, but for those of us just boarding the proverbial AI train, this information provides a necessary orientation.

Without bogging readers down in minutiae, Mollick's brief history of recent technological advances in AI gives readers a sense of how we got to the present day where something called a Large Language Model ("LLM") can power a bot that emulates human writing and thinking. With the reader focused on the relevant technology for today's AI, Mollick then explains how these LLMs operate, who created them, and how those creators built these systems. Specifically, he walks through the iterative "pretraining" and "fine-tuning" processes that LLMs go through, including a discussion of "tokens," "weights," and "Reinforcement Learning from Human Feedback." This discussion lays a strong foundation for the rest of the book and Mollick's overarching argument.

The Jagged Frontier

Mollick's mission in *Co-Intelligence* is to convince readers to use AI in their daily lives. This is because he needs their help mapping "the Jagged Frontier of AI."[9] Given that the universe of AI's potential capabilities is

7 *Id.* at 5.

8 *Id.* at 5–10.

9 *Id.* at 47.

so vast, Mollick wants people from a diversity of fields to test the technology and share any discovered strengths, weaknesses, possibilities, or limitations. Armed with the revelations from this kind of crowdsourcing, Mollick believes that AI researchers and developers can continue to improve AI, which in turn will lead to more innovation and, hopefully, societal benefits worldwide.

This grand vision is unsurprising, coming from a professor of innovation and entrepreneurship. But readers don't need to share Mollick's worldview to benefit from his book. Mollick gives readers a more straightforward charge: "try inviting AI to help you in everything you do, barring legal or ethical barriers."[10] Doing so, Mollick notes, could lead readers to enjoy the productivity gains, increased job satisfaction, and other career benefits that recent studies have seen with workers who use AI.[11]

And for readers who might still refuse to engage for fear of job security, Mollick also has a response. He concedes that no one can predict the effects of AI on the workforce and economy, and he acknowledges recent research that suggests that most jobs will overlap with AI's capabilities. But he reassures readers that, although this overlap will likely cause most jobs to change, it will not necessarily mean that AI will replace most jobs. Mollick thinks about jobs as "composed of bundles of tasks."[12] He predicts that AI will take over some tasks for every job, but he quips that workers may welcome offloading some of those tasks. And he believes that this reallocation of tasks will free workers up for more meaningful or important tasks.

Because the learning curve for working with AI can be frustratingly steep, Mollick doles out practical advice for using the new technology beginning in Chapter 3 and continuing throughout the rest of the book. A simple example is his recommendation that readers always plan to review and edit the AI's output before relying on or using it. But there are more detailed directives, too, such as his framework for determining whether and how to delegate a given task to AI.[13]

10 *Id.*

11 Mollick cites early AI research that showed that "[p]eople who use AI to do tasks enjoy work more and feel they are better able to use their talents and abilities." *Id.* at 153. And he cites recent studies involving participants writing documents they would typically prepare as part of their own jobs: "Participants who used ChatGPT saw a dramatic reduction in their time on tasks, slashing it by a whopping 37 percent. Not only did they save time, but the quality of their work also increased as judged by other humans." *Id.* at 111; *see also id.* at 126–27 (discussing similar studies Mollick is involved in with Boston Consulting Group).

12 *Id.* at 124–25.

13 *Id.* at 130–37.

A Pipeline of Humans in the Loop

Building on the foundation he laid in Chapter 1's discussion of how LLMs operate, in later chapters, Mollick emphasizes the potential dangers of an unchecked AI in the present day and the near future.[14] This isn't fearmongering to no end. He includes these warnings to lay the responsibility at the readers' feet. He urges them to become fluent with AI so that they can "learn to be the human in the loop."[15] By this he means that readers need to have enough working knowledge of AI to be able oversee it effectively, offering their own critical thinking skills, ethical considerations, and subject-matter expertise.

As Mollick points out in Chapter 8, most professional workers receive significant on-the-job training long after their formal education ends. He argues that AI puts this "hidden system of apprenticeship" in jeopardy.[16] As he sees it, working with humans can be emotional and inefficient. So, his argument continues, if AI now allows the boss to do certain tasks efficiently on their own, the boss is less likely to invest the energy and time in working with an in-training human. And since the boss holds some expertise in the profession, the boss's decision not to train the apprentice amounts to a decision to not share expertise. Over time and at scale, this creates a training gap that ultimately leads to fewer experts in society. And such a state of affairs would be deeply problematic because it would eliminate the very experts who should be the humans in the loop overseeing AI going forward.

AI in Legal Practice and Education

Thinking about which humans would remain in the loop in legal practice and education led to my own three sleepless nights. Legal practice is typically thought of as having a rich tradition of apprenticeship. But it's also a profession that places a premium on efficiency (though not at the expense of accuracy). That focus on efficiency makes law practice particularly vulnerable to an AI-caused training gap. Current experts from the bench and bar should pay close attention to Mollick's coverage of AI in business in Chapters 5, 6, and 8, and consider how they can best fortify their mentorship efforts to ensure that the next crop of lawyers and judges are properly trained.

14 Chapter 9 explores this topic further in imagining four scenarios for the future: "As Good as It Gets," "Slow Growth," "Exponential Growth," and "The Machine God." *Id.* at 193–210.

15 *Id.* at 52.

16 *Id.* at 178.

Mollick identifies a related AI threat in education. Effective on-the-job training depends on apprentices entering the workforce with some baseline education and requisite professional competencies. Mollick explains that AI's current capabilities and accessibility lead many students to believe that they no longer need to learn basic facts or amass basic skills in school. Mollick calls this the "paradox of knowledge acquisition in the age of AI."[17] He argues that acquiring foundational knowledge is more important than ever with the rise of AI because society needs a steady pipeline of expert humans who can oversee AI. Unless and until humans can acquire such expertise without traditional learning techniques of memorization, purposeful practice, and the like, educators have an important role to play.

Mollick's musings on teaching in the age of AI in Chapter 7 have a lot to offer legal educators. For starters, he doesn't put much stock in teaching prompt engineering.[18] Because current versions of AI can already figure out a user's intent, he predicts that in the very near term, that capability will be sufficiently improved to obviate the need for users to be good at prompting.

He also discourages educators from investing their time in designing low- or no-tech assignments and policies like in-class, handwritten assignments that prevent students from accessing AI. He views those as short-lived workarounds.

Instead, Molllick advocates for educators to focus on sharpening their evidenced-based teaching practices to ensure that AI does not prevent students from continuing to meet the educator's learning objectives for a given course or lesson. The recommendations in this part of the book will be familiar to many readers of this journal: Mollick emphasizes the utility of active learning in knowledge acquisition. And he sees flipped classrooms as a key feature of education's future. Delivering content to students (and their bots) as part of homework frees up precious class time to give students the opportunities for critical thinking, deliberate practice, collaborative problem-solving, and feedback.

He envisions educators crafting different categories of assignments and assessments depending on their goal—some will require AI use and others will forbid AI use, much like the variation seen in math classes with the use of calculators. In turn, educators will need to be transparent about their pedagogical choices for requiring or forbidding AI use. And

17 *Id.* at 181.

18 Mollick does mention basic tips for effective prompting throughout the book, though. These include providing context and constraints, using a chain-of-thought approach, and providing step-by-step instructions. *Id.* at 58, 170–71.

educators will need to have clear policies about etiquette and academic integrity surrounding AI.

Crowdsourcing

There's so much more to say about the contents of this book. It includes screenshots of interesting prompts to and responses from bots. It features a very accessible Notes chapter with citations to established scholarship in innovation and pedagogy as well as emerging (and not yet peer-reviewed) scholarship in AI. And it tees up but doesn't fully address many of the most heated debates in AI like the legality of LLMs' source text including copyrighted material, what it would mean for AIs to pretrain on their own content, whether we can assess sentience, how to regulate these technologies nationally and internationally, whether artificial superintelligence is possible, and much more.

Ultimately, the book leaves the reader wanting more. More information about AI. And more people to think through the future with. So, like Mollick, I find myself making an appeal for crowdsourcing. I am interested in having as many different people as possible from the bench, the bar, and the academy read Mollick's book so that together we can think about what this new co-intelligence means for legal practice, legal education, and (now I really sound like Mollick), *the world*.[19]

19 Mollick ends Part II of the book on a similar note:

> The thing about a widely applicable technology is that decisions about how it is used are not limited to a small group of people. Many people in organizations will play a role in shaping what AI means for their team, their customers, their students, their environment. But to make those choices matter, serious discussions need to start in many places, and soon. We can't wait for decisions to be made for us, and the world is advancing too fast to remain passive. We need to aim for eucatastrophe, less our inaction makes catastrophe inevitable.

Id. at 210.