

INSTANT NEW YORK TIMES BESTSELLER

ETHAN MOLLICK

Wharton School - University of Pennsylvania

Co-Intelligence



Living and Working with AI

These **InsightBites** were produced by: *InsightBites.io* Number of Text Pages: 13



15 min read



7 key insights

Contents

INTRO.....	3
Bite 1: The Concept of Co-Intelligence.....	4
Bite 2: Follow AI Working Principles.....	6
Bite 3: The Jagged Frontier Is A Thing	8
Bite 4: AI Has A Role In Creative Work.....	10
Bite 5: Human Oversight Is Needed.....	12
Bite 6: Human Values Before AI	13
Bite 7: The Future of Work In The Age of AI	15

INTRO

Co-Intelligence by **Ethan Mollick**, a Professor of Management at the Wharton School of the University of Pennsylvania, is a groundbreaking exploration of how artificial intelligence (AI) and human intelligence can collaborate to unlock new possibilities in creativity, decision-making, and work. Mollick, a leading expert in innovation and entrepreneurship, delves into how AI is not just a tool, but a partner that can enhance human capabilities. He illustrates how the fusion of AI's computational power with human judgment and emotional intelligence can lead to remarkable advancements in problem-solving and innovation.

The book outlines the principles and practices necessary to work effectively with AI, offering insights into its strengths and limitations. Mollick addresses the uneven capabilities of AI, showing how humans can navigate its "jagged frontier" through experimentation and oversight. By sharing examples from creative industries and professional fields, *Co-Intelligence* demonstrates how AI can amplify human creativity while reinforcing the importance of human oversight to ensure ethical and accurate outcomes.

Through real-world case studies and forward-looking insights, *Co-Intelligence* offers a practical guide for individuals and organizations seeking to harness the power of AI in their work. This summary distills Mollick's core ideas, providing valuable lessons for anyone looking to thrive in a future where AI and human intelligence work hand in hand.

We strongly encourage you to buy the full-length book.

This Book Belongs to the List:

AN INSTANT NEW YORK TIMES BESTSELLER (2024)

Find full-length book on: [Adlibris](#)

Bite 1: The Concept of Co-Intelligence

Ethan Mollick presents the idea of **co-intelligence**—a concept in which human intelligence works in partnership with AI to enhance cognitive capacities and improve decision-making and productivity. This approach redefines the traditional dynamic where AI is seen either as a mere tool or as a potential replacement for human labor. Instead, Mollick proposes that AI can function as a **co-worker**, complementing human creativity, intuition, and judgment, rather than competing with or replacing them.

At the heart of co-intelligence is the recognition that AI excels at certain tasks—like processing massive amounts of data, generating diverse ideas, or automating repetitive processes—but struggles with others, such as ethical decision-making or understanding emotional nuances. By integrating AI's strengths with human creativity and emotional intelligence, we can unlock unprecedented problem-solving abilities.

One of the key benefits of co-intelligence is its potential to **transform problem-solving and innovation**. Mollick provides a striking example from a business class at Wharton where GPT-4, an AI model, outperformed 200 students in generating product ideas. Not only did the AI generate a significantly higher volume of ideas in less time, but the quality of its suggestions also surpassed those of its human counterparts. Human judges, unaware of the AI's involvement, rated the AI-generated ideas as more marketable than those from students, highlighting the potential of AI to boost innovation by complementing human brainstorming.

However, Mollick is clear that AI's contributions must be **guided by human oversight**. While AI can generate vast amounts of data or provide creative solutions, it lacks the judgment needed to evaluate the moral, ethical, or emotional impact of its actions. This is where the human side of co-intelligence becomes crucial. Humans bring the capacity for ethical reasoning, empathy, and strategic thinking—qualities that AI, for all its computational power, cannot replicate. In this model, humans act as **stewards of AI**, ensuring that its outputs align with human values and goals.

A central theme in Mollick's work is that **AI's unpredictability**—what he calls the "Jagged Frontier"—is both a challenge and an opportunity. While AI can be remarkably creative in one instance, it may "hallucinate" or provide incorrect answers in another. This unpredictability requires humans to maintain a **continuous feedback loop** with AI, where they remain engaged and provide oversight to ensure accuracy and relevance in AI's outputs.

Mollick emphasizes that the concept of co-intelligence is not just theoretical but practical and actionable. By learning how to work with AI, individuals and organizations can take advantage of its strengths in areas like data analysis, creativity, and even decision-making, while maintaining the essential human qualities of judgment, ethics, and leadership. This collaboration not only enhances productivity but also empowers individuals to focus on higher-order tasks, such as strategy, relationship-building, and creative innovation.

In summary, co-intelligence offers a forward-thinking approach to integrating AI into daily workflows, making it a partner in creativity and innovation rather than just a mechanical assistant. By fostering a balanced collaboration between humans and AI, individuals and businesses can unlock new possibilities, improving decision-making, creativity, and problem-solving in a rapidly evolving technological landscape.



Bite 2: Follow AI Working Principles

Mollick introduces four essential principles for working effectively with AI. These guidelines are designed to help individuals and organizations leverage AI's capabilities while maintaining human oversight and avoiding over-reliance. Let's explore each principle in more detail:

1. Always Invite AI to the Table

Mollick's first principle advocates for integrating AI into every possible aspect of work. By involving AI in a wide variety of tasks—be it brainstorming, decision-making, or problem-solving—users familiarize themselves with AI's strengths and weaknesses. Mollick explains that the more we work alongside AI, the more we understand its potential for augmenting our work and where it might fall short. This principle underscores that AI should be treated as a general-purpose tool, adaptable to a multitude of tasks, rather than a one-size-fits-all solution.

This broad experimentation allows users to discover how AI can complement human efforts, especially in areas requiring data analysis or creative output. Mollick provides the example of using AI in complex, repetitive tasks like writing product descriptions, where AI can enhance speed and consistency without completely taking over the process.

2. Use AI as a Partner, Not Just a Tool

The second principle emphasizes that AI should not be treated like a traditional machine but as a collaborative partner. While many people tend to think of AI as just another software tool, Mollick argues that AI's true value comes when it works alongside humans, supporting them in tasks that require creativity, judgment, or nuanced decision-making.

Mollick compares this collaborative relationship to that of a "Centaur," where human and AI work together in tandem. The AI handles certain parts of a task, such as generating ideas or analyzing data, while the human manages higher-order tasks like refining ideas or making strategic decisions. This synergy allows for more efficient and innovative outcomes, as AI provides speed and data processing while humans contribute experience, ethics, and emotional intelligence.

3. Treat AI Like a Person (But Make Sure It Knows What Kind of Person It Is)

In this principle, Mollick humorously suggests treating AI as if it were a human partner—while remaining mindful of its limitations. This anthropomorphization helps users communicate more clearly with AI, giving it better instructions. By imagining AI as a collaborator, users can better prompt it to perform tasks more accurately and creatively. However, Mollick also warns that it's crucial to "tell AI what kind of person it is," meaning that users must clearly define AI's role in each task.

For example, in creative tasks, telling the AI to act like a famous artist or entrepreneur can yield outputs that align with that persona. However, humans must maintain control by setting ethical boundaries


and checking for errors, as AI doesn't "think" like a human and may sometimes generate inaccurate or misleading information.

4. Never Rely Solely on AI

The final principle reminds users not to become overly dependent on AI. Mollick emphasizes that AI, while powerful, is not infallible. It can generate "hallucinations"—wrong or nonsensical outputs that can mislead users if they aren't paying attention. Therefore, humans must always stay in the loop, reviewing AI's work and ensuring that its outputs are accurate and ethically sound.

In a study Mollick cites, recruiters who relied too heavily on AI to screen candidates missed key insights that human judgment could have provided, resulting in poorer hiring decisions. This demonstrates the danger of outsourcing too much responsibility to AI, highlighting the importance of maintaining a balance between AI-driven efficiency and human oversight.

Ethan Mollick's four principles guide users toward a balanced and effective relationship with AI. By integrating AI into tasks, treating it as a partner, setting clear boundaries, and maintaining human oversight, we can harness the power of AI without losing control or becoming over-reliant on it. This collaborative approach between human intelligence and AI offers a promising path forward, where innovation and creativity thrive through partnership.



Bite 3: The Jagged Frontier Is A Thing

The author introduces the concept of the **Jagged Frontier** to describe AI's uneven capabilities, illustrating how AI excels in certain areas while faltering in others. Mollick emphasizes that while AI can perform tasks related to creativity, analysis, and idea generation with remarkable skill, it struggles with tasks requiring basic math or logical consistency. This unpredictability creates what he refers to as a “jagged frontier,” a zone where AI's strengths and weaknesses co-exist, requiring humans to learn through experimentation and exploration to fully understand where AI is most effective.

At the core of this concept is the idea that AI's capacity is not uniform. In creative fields, like marketing or product ideation, AI can generate numerous ideas rapidly and with high quality, often surpassing human performance in these specific tasks. For example, in one of Mollick's studies at Boston Consulting Group (BCG), consultants using AI produced more creative, faster, and more analytical work compared to their counterparts without AI assistance. However, when AI was used for tasks it wasn't suited for—such as dealing with misleading or tricky statistical problems—the human consultants performed worse with AI than without it. This highlights the unpredictable nature of the Jagged Frontier and the necessity of understanding AI's limitations.

Mollick argues that humans must navigate this frontier by learning through trial and error. As individuals and organizations integrate AI into their work processes, it becomes essential to understand which tasks AI can handle efficiently and where human oversight is critical. The Jagged Frontier requires people to develop a nuanced understanding of AI's capacity, distinguishing between tasks where AI excels, tasks where it can assist but not replace human input, and tasks where AI's limitations are apparent.

This unpredictability of AI is not limited to technical or analytical tasks. Mollick explains that AI can hallucinate—offering creative solutions that are incorrect or completely fabricated. This is why human intervention is crucial. While AI might surprise with innovative connections, it also has a tendency to make glaring errors that a human would easily spot. Mollick shares a humorous anecdote where AI generated a fantastical solution for an unrelated task: when prompted for an idea involving fast food and a patent for a lava lamp, AI suggested a medieval-themed fast-food restaurant with decorative lava lamps for ambiance. Although entertaining, such ideas illustrate the unpredictable nature of AI and the need for careful human curation of its output.

Ultimately, Mollick emphasizes that the **Jagged Frontier** is not a static concept but one that will shift as AI technology evolves. AI's capabilities will continue to expand, but its weaknesses will remain, particularly in tasks requiring judgment, empathy, and moral reasoning—areas where humans excel. As AI improves in creative tasks, like generating business ideas or content, its struggle with consistency, fact-checking, and logical coherence will persist. Understanding and navigating this jagged frontier will remain a key challenge for organizations and individuals as they integrate AI into their work processes.

In conclusion, the **Jagged Frontier** serves as a metaphor for the unevenness of AI's abilities, calling for humans to be both cautious and experimental in their use of the technology. By recognizing the strengths and weaknesses of AI, and by working alongside it rather than fully relying on it, individuals and teams can achieve more innovative and efficient outcomes while avoiding the pitfalls of over-reliance. The concept emphasizes the need for balance—leveraging AI's creative and analytical power while maintaining human oversight to ensure accuracy, ethical decision-making, and thoughtful application of AI-generated insights.



Bite 4: AI Has A Role In Creative Work

The book explores the concept of AI in creative work, presenting AI not as a threat to human creativity, but as a powerful tool to amplify it. Mollick argues that while AI can generate a multitude of ideas and solutions rapidly, human involvement remains crucial in refining and selecting the best options. AI excels at providing creative assistance, but humans must guide the process by crafting thoughtful prompts and leveraging AI's strengths to push beyond conventional thinking.

Mollick's vision of AI as a "creativity amplifier" is particularly relevant in fields like marketing, design, and innovation, where generating a wide range of novel ideas is critical. For example, AI tools can help generate hundreds of ideas for new products, advertisements, or artistic concepts in a fraction of the time it would take a human team. In one case mentioned in the book, students using AI in an innovation class at Wharton were able to produce significantly better business ideas compared to previous classes without AI assistance. AI's ability to quickly suggest a wide variety of ideas enables human creators to focus on filtering, combining, and refining those suggestions to achieve truly innovative results.

Mollick also highlights that AI has the potential to democratize creativity. Many individuals feel they are not living up to their creative potential, with only a small percentage of people believing they are fully utilizing their abilities. AI offers a solution for those who struggle to generate ideas on their own, helping to unlock creative potential by offering suggestions, alternatives, and inspiration. For instance, Mollick shares his own experience with AI-driven art, where tools like MidJourney enabled him, a self-admitted non-artist, to create visually compelling artwork that resonated with thousands of people online.

Despite the positive aspects of AI in creative work, Mollick acknowledges the limitations and challenges. AI often produces "average" or crowd-pleasing results based on its training data, and it can struggle to generate true novelty without careful prompting. This is where human guidance becomes essential. Humans can direct AI to explore more unusual or high-variance ideas by framing prompts in a way that encourages exploration outside the norm. By doing so, AI can generate more original and unexpected combinations, which can then be honed into innovative solutions.

The key to successfully integrating AI into creative workflows, according to Mollick, is experimentation. AI's output is highly dependent on the quality and specificity of the prompts it receives. For example, asking an AI to generate marketing slogans for a new cheese shop might result in generic suggestions. However, by instructing the AI to act as an expert marketer and to focus on clever wordplay and varied ideas, the output becomes much more diverse and creative. This interaction highlights how human expertise and AI's computational power can complement each other.

One concern that Mollick addresses is the risk of AI diminishing human originality. As AI becomes more prevalent in creative fields, there is a temptation to rely too heavily on its initial output, which may lead to a homogenization of creative work. Mollick warns against over-reliance on AI-generated first drafts, which could stifle deeper exploration of alternative ideas and limit creative diversity. To counteract this, he advocates for a balance where AI serves as a starting point for inspiration, but humans remain responsible for pushing the boundaries of creativity.

Mollick also points out that AI has democratized creativity in a March 2024 Rotman University of Toronto Rotman InsightsHub blog post¹:

'In the most practical sense, we are now much less limited by ideas than ever before. Even people who don't consider themselves creative now have access to a machine that will generate innovative concepts that beat those of most humans (though not the most creative ones). Where previously, there were only a few people who had the ability to come up with good ideas, now there are many. This is an astonishing change in the landscape of human creativity, and one that likely makes execution, not raw creativity, a more distinguishing factor for future innovations.'

In conclusion, AI holds tremendous potential as a partner in creative work, enabling individuals to generate ideas more efficiently and unlocking creative potential in those who may not consider themselves particularly inventive. However, human oversight and experimentation are essential to ensure that AI's contributions lead to truly original and meaningful outcomes.



¹ <https://www-2.rotman.utoronto.ca/insightshub/creativity-innovation-business-design/ai-constraints>

Bite 5: Human Oversight Is Needed

Mollick emphasizes the vital role of **human oversight** in preventing AI "hallucinations," a term used to describe instances when AI produces incorrect or misleading information that seems plausible. This phenomenon is particularly prevalent in large language models (LLMs), such as GPT, which generate text based on patterns but do not store factual data or understand the world in the way humans do. These hallucinations are not just technical glitches; they arise from AI's inherent limitations in distinguishing truth from fiction and navigating incomplete or biased data.

Mollick explains that AI, while powerful in its computational capabilities, cannot reliably verify facts or ensure that the information it provides is accurate. As AI continues to develop and integrate into various industries, this issue becomes more pressing. He highlights cases where AI hallucinations have led to significant problems, such as when a lawyer unknowingly used fake court cases generated by an AI, leading to a fine and public embarrassment. Such examples underscore the importance of human intervention at critical stages to fact-check and ensure that AI's outputs align with reality.

The key to effective co-intelligence, as Mollick describes, is maintaining a "human-in-the-loop" approach. This means that humans must remain actively involved in AI processes to provide ethical guidance, critical thinking, and the necessary checks for accuracy. While AI can assist in generating content, solving problems, or streamlining tasks, humans are still needed to ensure that the results are credible and ethically sound. This collaborative loop prevents overreliance on AI, which could lead to errors, and allows humans to correct mistakes and refine AI outputs based on nuanced understanding that AI lacks.

Mollick goes further, discussing how hallucinations are not entirely negative. The same mechanism that causes hallucinations also enables AI to think creatively, making unexpected connections and solving problems in ways that humans may not have considered. In this sense, AI's potential for hallucination is part of what makes it a valuable partner in creative and exploratory tasks. However, this potential needs to be carefully managed through human oversight, ensuring that creativity does not come at the expense of factual accuracy.

Human oversight is, therefore, essential not only for preventing errors but also for ensuring accountability. As AI continues to evolve, its decisions will increasingly impact industries, legal systems, and everyday life. Humans must remain engaged, using their judgment to filter AI's outputs, guide its development, and adapt to new roles that require higher levels of cognitive and ethical responsibility.

In summary, Mollick's key message is that AI, while powerful, requires human oversight to be both effective and responsible. Without human involvement, AI's hallucinations can lead to costly mistakes, ethical dilemmas, and a loss of accountability. Therefore, humans must remain central to the AI process, balancing its strengths with careful supervision to ensure that AI remains a tool for good.

Bite 6: Human Values Before AI

Mollick highlights the importance of **aligning AI with human values**, recognizing that as AI becomes increasingly capable, ensuring that it serves human interests becomes more challenging. This alignment issue is not merely a technical problem for AI developers but requires a broader societal response. The alignment problem emphasizes that AI must be developed in a way that reflects human ethics, ensuring that AI technologies do not lead to harmful or unintended consequences.

Mollick points out that AI development is often driven by financial incentives, with companies focusing on rapid advancements. While governments and corporations must play a role in this process, the alignment of AI also depends on public involvement. Societies must set norms and ethical standards that guide AI's use and development. For example, AI's ability to act autonomously raises the question of who is responsible for its actions—whether AI companies, governments, or individuals. Without broad collaboration among these stakeholders, the risk of AI acting in ways misaligned with human interests grows.

The challenge of AI alignment extends beyond traditional regulation. Governments face difficulties keeping up with AI's rapid evolution, and global competition in AI development can push countries to prioritize innovation over ethical concerns. Even with regulatory efforts like the Biden administration's initial rules, there is still a need for more extensive and nuanced policies. These regulations, however, cannot address all of AI's potential risks. International cooperation is required, ensuring that AI advances are aligned with collective human values.

One key issue in AI alignment is AI's moral and ethical training. For instance, most large language models (LLMs) like GPT are trained on vast datasets scraped from the internet, which often include biased, outdated, or unethical content. This introduces the risk that AI will replicate harmful biases or produce morally questionable outputs. To prevent this, AI developers employ techniques like Reinforcement Learning with Human Feedback (RLHF), which uses human ratings to guide AI toward making ethical decisions. However, this process is not without cost, as human raters may face psychological distress from evaluating harmful or disturbing content produced by AI. Despite these efforts, no system is foolproof, and AIs can still produce incorrect or harmful outputs when manipulated through techniques like prompt injection.

Moreover, AI's training on biased datasets highlights another issue: the danger of algorithmic bias. Since AI reflects the data it is trained on, if the data is biased or lacks diversity, AI can perpetuate these biases. This can have real-world consequences, particularly in areas like hiring, law enforcement, and decision-making, where AI's outputs may reinforce existing societal inequalities. Mollick calls for transparency and accountability in how AI is trained, urging developers to consider the ethical implications of their data sources and models.

To achieve true alignment, Mollick argues, public education is crucial. The general public needs to understand AI's potential and its risks, allowing them to advocate for ethical AI use. AI is shaping our future in profound ways, from the workforce to healthcare, and its development must be informed by a wide range of voices to ensure it enhances human potential rather than diminishes it.

In summary, aligning AI with human values is not just a technical challenge but a societal one. It requires collaboration between governments, corporations, and the public to create ethical frameworks and regulations that ensure AI serves humanity's best interests.



Bite 7: The Future of Work In The Age of AI

The book emphasizes the **evolving future of work in the age of AI**. Mollick argues that while AI will take over certain routine or repetitive tasks, human expertise will remain vital in roles requiring judgment, creativity, and emotional intelligence. Rather than replacing human workers entirely, AI serves as a complementary tool, allowing people to shift their focus from mundane, repetitive tasks to higher-order, value-driven activities.

One of the most striking elements in this future is the creation of opportunities for humans to concentrate on areas that machines can't excel at—tasks that require personal intuition, deep understanding of context, and decision-making influenced by ethical considerations. Mollick illustrates how industries like medicine and education are already benefiting from AI augmenting human capabilities, such as assisting in complex diagnoses or supporting lesson planning. For instance, AI assistants can automate the more tedious aspects of academic research, freeing up time for human scholars to focus on innovation and critical thinking.

However, Mollick also cautions that this new reality will require significant changes in how society approaches work and employment. As AI increasingly handles analytical tasks, coding, and even creative production, workers will need to be retrained and re-skilled to maintain their relevance in this evolving landscape. Moreover, organizations must actively manage the integration of AI by ensuring workers are upskilled to collaborate effectively with AI systems, preventing a loss of human expertise and judgment.

Mollick's forecast sees jobs evolving rather than disappearing altogether. He highlights that AI is not simply automating tasks but also enabling new industries and roles that previously did not exist. For example, more specialists will be needed to manage and adapt AI systems, and creative industries will expand their horizons as AI continues to unlock new possibilities for innovation. The future of work, in Mollick's view, is a symbiotic partnership between humans and AI, with humans focusing on higher-value contributions while AI handles the tasks that are better suited to its capabilities.

To support this view, we can look at the projected numbers by experts. Sarah Choudhary, a Forbes Technology Council Member and CEO of Ice Innovations, said in a July 2024 Forbes Council post that²:

'A common fear is that AI will lead to mass unemployment. However, my experience shows otherwise. AI automates repetitive tasks, but it also augments human capabilities, enabling workers to focus on strategic and creative tasks. According to the World Economic Forum, by 2025, 85 million jobs may be displaced by machines, but even more—97 million—new jobs will be created. In our company, AI has streamlined operations, allowing our employees to take on more complex, fulfilling roles.'

² <https://www.forbes.com/councils/forbestechcouncil/2024/07/12/ai-and-the-future-of-work/#:~:text=AI%20automates%20repetitive%20tasks%2C%20but,new%20jobs%20will%20be%20created.>

To conclude, Mollick's vision urges businesses, governments, and individuals to prepare for a future where AI is a co-intelligence partner, not a competitor. The integration of AI is inevitable, but it can lead to a more efficient, creative, and fulfilling work environment if human workers are empowered to leverage AI's strengths while maintaining their unique roles as decision-makers and innovators.



LEGAL NOTICE:

*Legal Notice on Content Summarization for Commercial Use By **InsightBites.io**:*

Our **InsightBites** distill the main ideas and most important concepts from non-fiction books, but they are new, original works of their own, written in our Insight Bites style, and format. We work with publishing partners to select the most compelling ideas to **bite**, introducing customers to new books and authors they might otherwise miss.

Our commitment goes beyond offering quick summaries; we strive to respect the intellectual endeavors of all authors and publishers whose works we engage with. This legal notice serves to clarify our stance on copyright, the originality of our content, and our dedication to upholding the highest ethical and legal standards in our service. We urge all users, contributors, and interested parties to carefully read through this notice to understand our operational framework and values:

1. Original Content: All content under the "InsightBites" brand, including but not limited to summaries, interpretations, and analyses, is the original work of the InsightBites team. Our team undertakes extensive efforts to ensure that our materials are created from a fresh perspective, using original phrasing, structure, and presentation.

2. Transformative Nature: Insight Bites Content is not a mere reproduction of the source materials. Instead, it's a transformative work that provides a distinct and unique experience to our users, different from consuming the original content. Our purpose is to offer concise insights, breakdowns, and overviews, which are crafted independently, ensuring that they **don't replace** the original work **but rather complement** it or inspire further exploration.

3. No Infringement Intended: Insight Bites respects the intellectual property rights of authors, publishers, and all third parties. Our service operates within the parameters of both EU and US copyright laws, emphasizing the creation of original, derivative works rather than reproducing copyrighted materials. Our goal is to promote books and provide a new avenue for readers to engage with content, not to detract from the original works.

4. Feedback & Concerns: If you have concerns about any content on Insight Bites, please reach out to us immediately at sales@insightbites.com. We're committed to addressing all concerns promptly and ensuring that our platform remains compliant with all relevant copyright laws.

THE END

B