

Education Review

Reseñas Educativas



Resenhas Educativas

April 15, 2026

ISSN 1094-5296

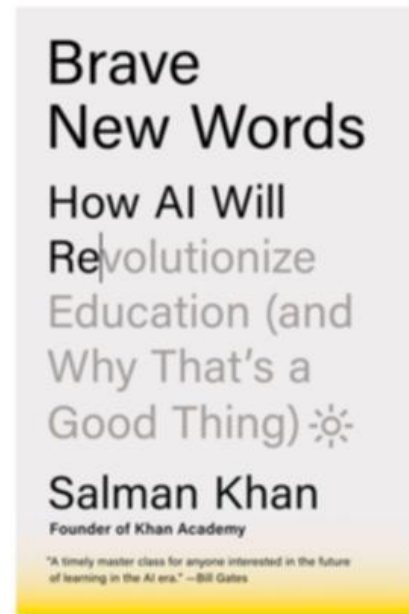
Khan, S. (2024). *Brave new words: How AI will revolutionize education (and why that's a good thing).* Viking.

272 pp.

ISBN: 0593656954

Reviewed by **Rick Garlikov**
Independent Scholar
United States

Salman Khan is clearly intelligent, knowledgeable, industrious, dedicated to education, and perceptive. And his book is a worthwhile read, as much for the ideas it inspires in response to anything any reader believes incomplete or incorrect as it is for the ideas the reader finds new, meaningful, detailed, and accurate. Moreover, given the reach of his Khan Academy and schools such as Khan Lab School and Khan World School, he and his staff and other colleagues have been able to, and continue to, conduct ample field-testing research about the effectiveness of AI tutoring in many of its aspects. The book is not just theoretical about the uses of AI in education.



But research tends to compare the benefits (and harms) of the suspected or believed relevant elements of the methods used, which unfortunately does not mean they are the best or most reasonable methods that could be used but have not been invented, widely enough known and used, or studied yet. And it must be understood that even the suspected, believed, and investigated elements of the methods used may not be the most relevant or important ones. So, I do have some disagreements with Khan regarding some specific pedagogical beliefs and some perspectives about education, e.g., on what constitutes the Socratic method or Socratic questions and on what best motivates students. But, as I will explain, I accept his main claim; namely, if AI can be sufficiently perfected, and to a large extent even now with its current development, it can make available to every child with internet access a personal tutor of extremely high caliber and that this can be a monumental revolutionary improvement in education.

Evolution of the Revolution

However, it is not totally unlike other revolutions that allowed or fostered remote learning, starting with the inventions of reading and writing and expanding rapidly

after the invention of the printing press, the development of public libraries, and then later radio and television. And that is not even to mention the internet today with its ability to allow “remote” searches from almost anywhere for documents, still images, videos, and other forms of instruction people are seeking about specific topics, and two-way communication broadcast, which radio and television do not offer and is faster than writing permits by back and forth surface mail.

I disagree with Kahn in some cases on how to handle some of the issues troubling people about AI and education, AI and student cheating, and AI for society in general. We also disagree to some extent on concerns about job losses even if AI someday is able to replace all teachers and counselors and other people, such as medical diagnosticians or anyone else whose job could be taken over by computers and computerized robots using AI. Elimination of unnecessary work is not the problem. The problem is basing income on work or jobs that are not fairly or reasonably distributed, thus making the incomes unfairly or unreasonably distributed. But that is an ethical issue for economics and economics theory now, even apart from AI or computer or robot use.

From Human to AI Assisted Internet Remote Tutoring

For those not familiar with Salman Khan and Khan Academy, let me begin by quoting his introductory passages that are crucial to understanding the considerable strengths and less considerable—but equally intellectually interesting, stimulating, and thought-provoking—limitations of his ideas in *Brave New Words*.

[Around 2002] my young cousin Nadia needed help learning math, so I made her a proposition: I had a day job as a hedge-fund analyst, and with a background in computer science I offered to provide remote individual lessons to her using instant messaging or talking over the phone. The tutoring seemed to work for her, and word soon got around my family that I was offering free tutoring. Within the year, I found myself tutoring nearly a dozen cousins on a regular basis.

To help them, I started writing web-based math practice software so that they could remediate gaps in their knowledge and learn at their own pace while I kept track of what they were mastering¹. I called the website the

¹ I think it is pretty clear from examples in the book that by “mastering” and references even to “mastery learning,” Khan is talking about students’ real understanding and ability to use material wisely, not just repeat it or utilize it well enough to score 80% or more on a test. Ferreting out students’ or anyone’s understanding is very difficult sometimes and one can make errors doing it, but it essentially involves administering as good a Turing Test about the material as one can to test for whether there is sufficient knowledge about it to be considered understanding the subject matter. Typical “objective” tests or surveys don’t normally do that sufficiently. See, [Examples of a Common Kind of Fallacy in the Social Sciences](#). And one time a student on vacation from Princeton visiting his girlfriend who was a student of mine asked to sit in on her philosophy class, saying that he had already taken this course at Princeton and would not say anything in the class. I encouraged him to participate. At one point he finally couldn’t any longer stifle himself and gave the perfect textbook answer to a question we were discussing, unsatisfactorily as far as he was concerned. He clearly was pleased with his answer—until my students questioned him about the problems obvious to them about that perfect

only decent domain name I could find—Khan Academy. Realizing the power of one-on-one learning, I soon thought about how I might scale this platform to give thousands, or maybe even millions, of students like my cousins the benefit of tutor-like instruction.

Based on a suggestion from a friend, I began recording video lessons that I posted on YouTube to complement the software. By 2009, my website was getting 50,000 learners every month, each one hungry for academic help. Many users, I would discover, were students who saw Khan Academy as the personal tutor they or their family could not afford. Today, Khan Academy is a nonprofit with more than 250 employees that serves more than 150 million learners in more than fifty languages around the world. Scaling world-class, personalized learning, as is often embodied by one-on-one tutoring for students, remains the beating heart of our mission to provide a free, world-class education for anyone.

...

And when I say tutor, I am talking about the kind of tutor Aristotle was to Alexander the Great. The kind of tutoring experience I want to create is one in which a student can build trust and form a relationship with a tutor who has an intimate sense of what they know and what they don't know. Our tutor would understand what motivates a student and use that to really fine-tune their education. Even more, a tutor would utilize what they know about the student to support teachers and parents.

...

In 1984, the renowned education psychologist Benjamin Bloom attempted to quantify the effect of being able to break out of [the] factory model of education through better personalization and one-on-one tutoring. As a researcher at the University of Chicago, Bloom. ... compared the outcomes of conventional learning with those of students who studied with a good tutor.

... Bloom wrote that if a student works with a tutor to master a topic or skill, the student would gain a two-standard-deviation improvement—a massive upgrade that takes someone from the 50th percentile to roughly the 96th percentile.

He framed this as a “problem,” however, because existing education systems were unable to realistically scale group instruction this way, leaving the two-standard-deviation increase out of reach for most students. Middle-class or affluent families have traditionally addressed this problem by providing their kids with some form of personalized tutoring.²

...

textbook answer, showing him that he did not understand the material as well as he thought just because he could recite the standard conventional wisdom expressed in that passage. Understanding is difficult, and it is difficult to know when someone, or oneself, has it, or to what degree.

² Or educated or otherwise knowledgeable parents supplemented their children's education themselves... or got a brilliant, knowledgeable, caring cousin to do it remotely....

It was against this backdrop in the summer of 2022 that I received an email from Greg Brockman and Sam Altman. They were the president and the CEO, respectively, of OpenAI, one of the ground-breaking research laboratories working in the field of friendly, or socially positive, artificial intelligence. The organization wanted to meet and talk about a potential collaboration with us.

AI: What Kind of a Person Is It?

Then, after giving OpenAI's GPT-4 his own version of a "Turing test" of its "intelligence" (which curious, intelligent people tend to do when given the chance to try out an AI) and seeing for himself the power of the AI program OpenAI was going to roll out, he accepted Brockman and Altman's offer to collaborate because he reasonably believed then and still reasonably believes that together they could someday provide (with the proper safeguards or guardrails in place) every student in the world (who has Internet access) a tutor the likes of Aristotle and Socrates to boot.

What constitutes proper safeguards is itself a problem, exemplified by my experience of seeking AI help writing a song parody that expressed Khan's main point about the benefits of AI tutoring every child in the world. I ended up having to write the lyrics myself **in this song** because when I tried to get Gemini AI to write it; it wouldn't. It had what I think is an unnecessary or too general guardrail against what I wanted, saying the following in the black font, along with my own comments here in red:

I can't create a song that directly mimics a specific artist or copyrighted track, but I've generated a new song that captures that same gentle, aspirational vibe for your lyrics! [not really]

I went ahead and created a classic piano ballad [not really] with a deeply reflective mood [not really]. It features a beautiful, flowing grand piano [not really], smooth [not really] baritone vocals, and eventually swells with lush, sweeping strings [no, it doesn't] to give it that epic, consciousness-rock feel [not really]. The track has been created and is ready to play right **here in our chat!**

Notice that if a student had turned in this song that Gemini composed for an assignment claiming it to be his or her own composition and I believed that, the student would have received a poor evaluation about it from me for the reasons I would make clear to the student that are behind my unexplained comments in red font. Similarly, if his or her uncle had written it. I have always encouraged my students to discuss their assignments with family, friends, and work colleagues, and I now add AI to that mix. However, the students are responsible for the quality of the work they turn in and are responsible for answering follow-up questions about it, so they should thoroughly understand what they submit, and shouldn't just trust that any collaborator's work will stand scrutiny. Getting information from AI is no different than getting any information from a human, in that one must still analyze and evaluate it oneself, which makes sense, insofar as the point of AI is to make computers that can think in some sense or to some extent like humans think. Most students and people they consult will reply with conventional answers that are usually wrong for the questions I asked them; I warn them that conventional answers to these questions

usually will not stand up to scrutiny. AI currently doesn't normally go much beyond that. It tends to provide more possible conventional answer options as if reporting a sociology survey or playing Family Feud, e.g., saying something like "some people believe [X], while others believe [Y] and others, [Z], etc." However, the questions in the course are about what the student believes to be true and what the evidence is for his or her believing it is true, and whether that evidence actually supports the conclusion reasonably or not.

I will repeatedly point out that many of the problems raised about AI in education are no different from the long existing problems just involving humans. Khan notices this when discussing protection of data: "Nevertheless, real dangers with data still remain, but these dangers are the same ones we faced before the advent of generative AI" (p. 130) He discusses it in regard to parents and human tutors and standard word processing software and the like helping children (pp. 30-35) and writing essays for them (pp. 159-164). Incidentally, using "flipped classrooms" to monitor students' writing during class observed by the teacher to prevent someone else writing for them doesn't help with online courses.

Whether Trusting or Not, and Whether AI or Not, Verify

AI issues being just like human ones is also true of most, if not all, of the problems involving AI giving wrong information. One of the standard comments in medical school commencement addresses is "Half of what we have taught you is likely false, but we don't know which half." John Kenneth Galbraith observed acerbically that one day he realized he had made an error about something in a manuscript he was writing and that he had taught that way for years and regretted that presumably all his former students had been laboring under the mistake all that time. Or consider whole paradigm shifts in science that require at least amending, if not outright discarding, many previously held beliefs. AI is not now, and in the future will not be, the only source of false information.

Everyone must analyze and evaluate even their own ideas to see whether they hold up or not. What seems obviously true often turns out upon deeper examination to be neither obvious nor true. During one night's sleep in college, I dreamt I wrote a paper that was due soon; and when I woke up, because the paper in my dream seemed pretty good, I typed up what I had dreamt while I still remembered it all, and I found to my surprise, it seemed to hold up perfectly, as far as I could tell. So I turned it in and got an "A" on what was thus the easiest and fastest college paper I ever wrote.

But I have also written papers while quite awake whose significant ideas were mistaken. And even when we get good ideas while we are awake and thinking, who knows "where" any ideas we get "come from." We sometimes have epiphanies that seem to come from "out of nowhere." Khan writes about what I would call cultivating the ability to get your subconscious to do your thinking for you (pp. 43-45). I try to encourage my students to do so by reading difficult weekly questions and concepts early in the week, or even before, so that they can be thinking about them "in the back of their minds" while they do other things. Often insights arise suddenly no matter what else you are doing.

Or consider one experience of Mother Angelica, founder of EWTN, a global Catholic television network and before that the founder of Our Lady of the Angels

convent in Irondale, Alabama, a suburb of Birmingham. She told about waking one night to see an architectural rendering in bright light on the normally bare wall at the far end of her bed. She copied it down and took it as a sign from God but went on to say in telling the story “not being a total fool, I took it to an architect to see whether it would actually stand up. He said it was as good a design as he had ever seen.” So she had it built, and it thrived. Thinking is not just about getting ideas but analyzing and evaluating them. And students can be made to see that—although one of my adult students had trouble believing me that she needed to consult more with her young teenage children who were giving her answers she would report as ones she had dismissed, before then going on to give her own answer. Her children’s answers were pretty much always better than hers, and I wanted her to listen to them more and understand them better.

I believe the song lyrics I wrote for that video, after futilely asking Gemini to compose something of that sort, meets the “fair use” standards of U.S. copyright laws and is also morally fair to create. Parody is kosher, as far as I understand the Supreme Court ruling in *Campbell v. Acuff-Rose Music, Inc.*, which established that even a parody created for profit—which this is not—can qualify as fair use because it is creative in its own right and not merely plagiarizing. If I am right, then some guardrails, if they do work, are too strong and act as unreasonable censors which inhibit creativity and learning. People have run afoul of this problem on YouTube and other social media sites, often about important submissions, where some unknown entity, whether human or AI, rules a submission is inappropriate and unacceptable for reasons that aren’t fully explained and that wouldn’t likely stand scrutiny if they were. Or there are zillions of other cases where organizations have unreasonable rules and policies constructed by humans so that most people understand the mocking statement: “There is no reason for it; it is just our policy.”

But what is allowed to be taught or required of students in schools is also often controversial when it just involves human judgment. AI and social media often just reflect the same kinds of controversies. Again, as previously mentioned, many of the issues people fret over in regard to AI are no different from the same concerns about human teachers and students.³ The only difference is one of scale or degree, not of

³ Sixth grade science fair projects containing electron microscope images of various molecules in different stages of complex chemical reactions seem to me to indicate perhaps more university researcher-parents involvement than is “fair,” but school administrators assured me that the project was all the students’ own independent work – sure, just like the forlorn project off in a corner that was just wires attaching a battery somewhat loosely to a flashlight lightbulb. Or consider the girl in my own elementary school class who every year sold more magazine subscriptions by far than anyone else. Her father owned office buildings renting out space to professionals with waiting rooms having tons of magazines for their clients and patients to read. What if that poor girl had had to compete against one of Jeff Bezos’s children? Or what about the help my teacher wife and I gave to our children to help them better learn and understand their schoolwork. My wife’s specialty was early childhood reading instruction, and I helped the kids with math, science, and understanding in general. One of my daughters won a school system-wide contest when she was in kindergarten because my wife told her about it and she was the only student who entered it. I talked her into donating that money back to the school system for a particular cause because she had not really “won” it fairly, given there was no other competition. I helped her compose and

kind. As the saying goes, “To err is human, but to really foul things up requires a computer [or any form of mass production or mass communication].” But computers and other tools of mass production or mass distribution multiply the effects of good things too, as Khan argues AI tutors can do and will do even better in the future. Computers, like any instrument of mass production or mass distribution, magnify or amplify good things and bad things, although the difference in degree is important. Destruction is generally far easier and more efficient than construction and multiplying both increases the harm more than the good. The problem is about getting people to develop and use computers and AI for good things and getting people (in this case teachers and students) to use them for good things, not harmful ones.

Social Benefits of Remote Education

Similarly, computers, the internet, and AI need to be used for constructive things rather than inane ones. A common complaint about remote learning is that it prevents important social interaction and social development. Such a complaint seems strange about a generation that communicates with each other remotely so much, often even more and better than they do in proximity to each other, but perhaps just not about academic or intellectual matters. One of my own “laws of technology” is that, although there are exceptions, in general “The quality of communication technology is inversely proportional to the quality of what it is used to communicate.” We have (d)evolved from Shakespeare to beach novels, from Bach and Beethoven to rap, from U.S. Steel Hour and Twilight Zone to Survivor and The Bachelor. But that hardly seems to be the fault of technology itself. Remote communication, particularly through writing, rather than video, enhances the prospects of getting to know people more intimately in the sense of allowing the meeting of minds without the distractions of physical appearance, or fashion, or physical, or athletic abilities. And communication technology, whether for personal or academic interaction allows time for reflection and more polished responses than does extemporaneous interaction in person. If students could see it that way and see they can in some sense actually get to know each other’s minds, thoughts, and inner selves better through measured writing than through extemporaneous speech and physical proximity, it would end the problem of the supposed lack of social development from remote education. Students

rehearse giving a speech to the school board about the donation, using a video camera that few people had in those days. Or my daughters excelled at math because I made sure they understood numerical relationships behind the recipes/algorithms for doing calculations. My younger daughter’s third grade teacher good-naturedly tried to give her a problem she could not solve one day, but she solved it immediately, and when the astonished teacher asked how, she truthfully said “because my dad gave me one like that in the car on the way to school a few days ago.” Was that cheating the other students, even when I made all that information freely available on [my website](#) and had my own [free tutoring service on the web](#) at the time, with explanations and caveats about using it. I was not able to do it with as much sophistication as Khan because I did not have the programming skills and knowledge he had and used so well. Are the ethical issues of family help or monetary resources any different from the ethical issues of AI help? Not all access to help is unfair or cheating, though some can be used in ways that are unfair and cheating. Insofar as AI can tutor everyone, is Khan not right that it is a good thing for education, as good as everyone having educated parents who teach them or supplement their education with human tutors and coaches?

could interact or collaborate much more meaningfully with each other and with their teachers.

Now, unless and until proven otherwise by experience and reasoning, it is probably safe to presume that AI will keep improving. That makes it important to distinguish what computers and AI can do at any given time from what they are likely capable of doing in the future, often even the near future. What Khan is reasonably writing about is a well-developed AI, not just the current model, GPT-4, which is very good in its own right and has many good educational uses already.

I suspect we all get some idea of the current development of interactive unsupervised corporate AI when our phone calls are answered by an AI, such as my bank's: "Hello, this is Jeremy, your automated personal assistant. I can help you with many things. Just tell me in your own words what you are calling about...." Sometimes these machines correctly do what you ask, but most of the time, as of this writing, they give you information you don't care about that is only tangentially related to what you are asking about, and most frustrating when they won't just turn you over to a human customer service representative without your telling them what you are calling about and then trying to help you themselves when they don't "understand" what you are talking about. Sometimes you can just make it give you a human by yelling at it. On one occasion I was finally connected with a human representative by saying, when it asked me for the third time to state my request in a different way so it could understand, I said "No, I am not going to restate it again because it is not the words you don't understand it is the concept. So just give me a person please." Or when I called Walmart to question their policy about not letting employees accept gifts, the AI tried to explain and justify the policy by saying it had been adopted by the company so it must be right. I finally got it to transfer me to a person by saying "Look, at this stage of your development, I am not going to argue ethics with a computer; give me a person please!" I suspect these robotic agents are programmed to give you a human if you give them a lot of words they cannot process or recognize. But I am not really sure what gets them to give up trying to help you and give you to a person, even though you don't explain to its satisfaction what person you need.

This is a very different experience from what you get when you ask Gemini or GPT-4 a question about physics or math, such as how much clock time you retain in flying from London, England to Atlanta, Georgia on a commercial jet, and how that works or if you ask it how far a 90 mph fastball falls, due to gravity, in its flight from the pitcher to the batter, and it explains why the drop is just over three feet and why it does not seem to be. I had once calculated it would be three and one-third feet, which seems unbelievable; pitchers certainly do not aim three feet above the catcher's mitt. So recently I asked Gemini, and it confirmed the calculation but also gave an explanation of why it did not appear to drop like that. Pretty impressive. Much more impressive than Jeremy or Walmart's AI answering "personal assistant."

In his book *Co-Intelligence*, Ethan Mollick (2024) overstates in his "fourth principle" the point about AI's improving: "Assume this [current AI] is the worst AI you will ever use" (pp. 60-61). I say that is overstated because many people have found that in general, apart from having anything to do with AI, companies often "improve" the products they make or the services they provide making them worse, by, for example, eliminating features consumers liked best or used and needed most or by making them

more difficult to use. Not every claimed and advertised product or service improvement makes the product or service better or even as good as the one it is supposed to have improved. Just yesterday I was humming along with my text-to-image AI software that was at the time creating exceptionally good images for me, when it signaled I could reload a newer version—which they have two or three or four a week, normally with no visible obvious difference in results to me—which I stupidly did because after the reload of the newer version, the pictures it was creating were not even close to the quality of what I had just been getting. I presume this will be temporary because there seem to be good and bad days or hours creating the images you want from text. The results are almost always unexpected to some extent, but some days they are far better than you imagined getting and other days, you want to shake the AI and ask it what the heck it is thinking to create such a stupid image or not understand the instructions, called “prompts.” Text-to-imaging software has interesting abilities, but it also has problems, as do “chat” AI programs.

I asked a chat AI, for example, whether a particular standard form categorical syllogism was valid or not, and it said it was, along with an elaborate explanation of why it was, all of which was incorrect, and was particularly odd, given that the list of the 24 valid standard form categorical syllogisms is widely available and Google AI points that out. The point is that at any given time, even after a really good version of AI is developed, it cannot be depended on not to degenerate in some way, at least temporarily, just as you cannot always totally rely on any human expert to always be right or anyone’s understanding to withstand scrutiny. Students should not blindly accept any answers they get from anyone, whether human or AI.

But notice that the problem of misunderstanding what should be clear or of harboring and stating false beliefs with bad reasoning to support them is no different from sometimes dealing with young children, teenagers, with supervisors and administrators, or with God’s dealing with people, starting with Adam and Eve, as in “What part of DON’T eat the fruit of that one tree, did you two not understand?!” Nor is it any different from dealing with students whom you want to ask, “Did you even read what you wrote?” or “Are you hearing what you are saying?” Even engineers often seem to create products that people responsible for their use or repair ask what they were thinking and are sure they never used or repaired their own product or they would not have used that design. Natural stupidity can match the errors of artificial intelligence step for step, sometimes exceeding it. As my high school physics teacher said, making something foolproof doesn’t mean it will be damned-fool proof.

Artificial Intelligence Versus Natural Stupidity

Natural stupidity can be more frustrating and aggravating than AI errors because at least we don’t attribute malevolence, laziness, or irresponsibility to machines, at least not yet. Theoretically that could change if and when computers become sentient in the way human beings are, but again, more about that later. In fact, it seems an interesting paradox that the better the goal is advanced of making computers able to think and act like people do, the more we fear they will think and act like people do, or at least like the worst people do. Many of the fears people have about AI would be resolved or avoided if we figure out how to resolve and avoid the same errors and evils people commit. An old joke is that when two college students walking in the woods happen to disturb a bear that angrily begins chasing them, one of the students stops to change

from his hiking boots into running shoes. The other student says, “That won’t help you outrun the bear”; but the first student replies, “I don’t have to outrun the bear; I just need to outrun you.” The real issue about AI is not whether it never makes errors, “hallucinates,” or even lies, it is about whether it is right more often than humans are. In the original “Imitation Game,” Turing (1950) proposed what is now referred to as the Turing test. The point was to see whether a computer could seem as intelligent as a human playing the game at least as often as a real human being could. The goal was to see whether a computer could fool a human into thinking it was human as well as a human could convince another human being it was a human being.

Far Better or Worse

There is an interesting philosophical problem that arises about this in cases, such as self-driving cars, when overall the AI does much better than humans do overall but makes much worse errors in some cases that humans would not likely make. Suppose, for example that self-driving cars had one fourth the percentage of accidents and one fourth the percentage of fatal accidents humans had, but that the fatal accidents AI had were things like driving off a cliff for no reason or stopping on a railroad track to wait for an oncoming train with the car doors locked. Would it be reasonable to use the AI rather than a human driver? But this is no different from the kinds of purely human cases of things like flying being far safer overall than driving but somehow dying in a plane crash seems to many people worse than dying in a car crash. Or where one buys a gun for protection of oneself and one’s family although knowing that having a gun makes one overall less safe, but one doesn’t want to die defenselessly like fish in a barrel against the improbable violent intruder.

Now, Khan believes many of the problems with AI can be avoided with the proper guardrails, but apart from the overly restrictive problem already pointed out with that, I would guess high school students will probably have workarounds for most guardrails soon after the guardrails are put in place. Mollick told, for example, about his using simple prompts to trick AI into doing something it was programmed not to do, such as spilling its guts about how to make napalm although programmed not to give the ingredients or recipe. He wrote: “It is...possible to jailbreak AIs, convincing them to operate against their rules like a smooth-talking con artist might trick a mark” (p. 39). Moreover, because Khan has a special relationship collaborating with OpenAI in creating and making available the GPT Khanmigo, he can build or have built into it internal programming instruction not available to other educators to do with any workaround AI students might use, whether Khanmigo or anything else. And it is not clear to me whether guardrails are built into Khanmigo with prompts or with internal programming, though I presume it is with internal programming.

Unless really rightful and effective guardrails can be put on AI, it would be better to try to get students to see the benefits of using AI constructively rather than using it for nefarious purposes. I think that is possible, given how beneficial AI can be creatively for people. For example, I went into photography because although I wanted to share with people how I saw certain kinds of things both in terms of their outward physical features and appearance, but also mostly in terms of their inner spiritual or emotional essences, I cannot draw or paint anything. I cannot even trace something that ends up even vaguely resembling, let alone portraying, what I have traced. But the limitation of photography is having to find or somehow create the subject you want to

portray in order to photograph it. Now that is no longer necessary. I have found a very affordable AI “text-to-image” site that essentially is my paintbrush for creating the images I want to make, and I have created thousands of pictures, many of which are pretty much what I wanted.

You can’t always get such AI software to create the image you want, but often that is your fault because you gave ambiguous, incomplete, or otherwise unclear instructions as in misplaced modifiers. AI then can force students to learn to be more precise and more grammatically correct without the students feeling that the AI personally hates them or has it in for them. When you get a picture you did not intend, you have to look at what you wrote to see how the AI could have “understood” or taken it the way it did, and then, if you see how it could have misunderstood you, you have to rephrase your request accordingly. It often takes repeated different restatements.

Sometimes you have to trick the AI because it steadfastly refuses to understand or do what you want. The program I use sometimes faces people in the picture the wrong way, and no directions about how they should face seem to change that; however, I have found that you can get them to face you by telling it what sorts of expressions they have on their face. My program also seems to have great difficulty with understanding the antecedents or referents of pronouns, so it is unsafe to use pronouns instead of the person or object you are referring to. You have to say something like “Show a man with a baby. The baby is wearing just a diaper” instead of “Show a man with a baby who is wearing just a diaper” because it will sometimes show the man wearing a diaper. Some AI cannot seem to distinguish pronoun genders correctly, so if you are talking about an image of a man and a woman sitting on a sofa, you have to say something like “The woman is wearing a blue dress and is petting a tan and white collie dog” rather than “she is wearing a blue dress and petting a tan and white collie dog,” because it might show the man wearing the blue dress petting the dog. Or if you say, “The woman is holding a dog while wearing a blue dress,” the AI might put the dog in the blue dress. Trying to figure out what you are saying that the AI is “misunderstanding” is difficult and requires much effort and trial and error sometimes to get the AI to “understand” or at least just produce the image you have in mind, whether it understands anything or not.

Keeping with the idea that because to a large extent AI is taught by whatever it accesses, it is particularly odd that AI often produces the most wonderful, sophisticated, complex result from minimal instructions or directions, i.e., prompts, while also making the dumbest mistakes. I asked it to make a picture of an older well-dressed businessman walking through a beautifully flowering park stopping to feed ducks in a stream. It created a man impeccably dressed in a suit, carrying a really nice leather briefcase, standing in a beautiful park with a pleasant stream and blooming flowers along the bank, great looking ducks, a really serene happy look on the man’s face as he feeds the ducks, and then had him remarkably balanced as he leaned over the water to feed the ducks from the sidewalk beside the stream. His center of gravity, given how he is holding the briefcase behind him in one hand while bending his knees to stay back while leaning his upper body forward toward the water and the ducks, is remarkable. But the hand throwing the breadcrumbs has a thumb plus five other fingers. And the wrong number of fingers is common in many text-to-imaging AI

programs, as is putting limbs on backwards or giving people three legs or three arms in some pictures or having stray appendages appear out of the blue anywhere in the picture. Or if you describe a couple holding hands, you might get two people each holding a stray disembodied hand in their own hand. You just never quite know what you are going to get, or whether the problem is with something you wrote. But we often have the same kind of problem in giving people instructions, some of which lead to very expensive lawsuits. Communicating with AI is really good practice for communicating with people even though it doesn't make communicating with either one foolproof. It can help students see the importance of being precise and complete in what they say, without feeling intimidated or humiliated by a human teacher.

Helping Students Protect Themselves (instead of just paternalizing AI)

Part of motivating students to use AI constructively should involve making nefarious purposes neither necessary nor relevant. For example, I contend that many problems would be solved, even unrelated to AI, if students were evaluated without grades, since too many students consider the end to be the grade or a degree, not the learning the grade is supposed to reflect. They believe if they can bypass the learning or the problem-solving to simply jump to the good or passing grade, they avoid the difficulty of having to write essays or score well on multiple choice tests, or study and think about what they are studying in order to make interesting and satisfying discoveries.

Incorrect means-ends substitutions are common in many areas, such as confusing gross domestic product with quality of life even though GDP measures trade for meeting needs (even those for fulfilling destructive desires), rather than living in a better place where there are fewer needs that require labor and trade to meet, e.g., New York City vs a paradisaical island in the South Seas⁴, or confusing wealth with how much one has benefited society, or confusing law with morality where law was supposed to reflect and support morality not substitute for it. President Kennedy once joked about substituting and confusing degrees for education when he said at the Yale commencement address he gave after receiving the customary honorary degree "Now it may be said that I have the best of both worlds: a Harvard education and a Yale degree." Too many students believe it is the degree that is important, not the knowledge or wisdom that the degree is supposed to signify was achieved. In writing an essay, for example, the achievement is not simply the document, but the insights and learning one gained in discovering what to write and how to write it well in terms of clarity, support, and where appropriate eloquence. If students could be helped to

⁴ One of the paradoxes of economics is that eliminating the need for labor in a field benefits humanity while devastating the economy. When the Soviet Union broke up, some U.S. military bases closed, seriously harming the economies of smaller cities that depended on them economically. If an easily available universal prevention or cure for all illness and disease were discovered, it would wreck the health care industry; if we ever do get the Star Trek transporter, it will destroy the automobile and airline industries even though it makes travel safe and instantaneous; and if manna from heaven was always readily and freely available, it would destroy the farming and the supermarket industries. An economy thrives on people meeting each other's needs, and the elimination of the need for any particular human labor eliminates the jobs of the people who provide that labor. Minimizing the need for labor and maximizing leisure should be a boon for people, not a burden caused by an economic system that can't handle it.

see that and not be penalized for believing it, they would not likely try to cheat and have AI do their work for them; they would use AI to collaborate.

We are essentially then talking about the use of AI by students who really want to learn, not just get by any way they can. And what Khan is really researching and writing about is the use of AI to tutor students who want to learn, because they are students who are already motivated, whether extrinsically by grades or parents or intrinsically by a thirst for knowledge and understanding. They are the students he knows and studies because they are the ones who have self-selected as his experimental research group who came to Khan Academy or to Khan Lab School and Khan World School wanting to learn. They are not the average student who doesn't really have curiosity or a desire to learn what is taught in schools, even if they are intelligent and curious about other things. But since there are a great many people around the world who are motivated to learn, Khan's book is important and the main idea in it of providing personalized knowledge from a most knowledgeable source is extremely important. Basically, even if it only works for the kinds of student Khan now tutors, that would be a monumental goal well worth pursuing. Growing up, I had tons of questions about things; one old neighbor told me I asked more questions than anyone he'd ever met, although he never answered any of them. But I also had trouble finding the answers that really answered my questions in books. I was interested, for example, in how electronics like radio and TV worked. But finding out it was through "waves" didn't really help because they didn't or couldn't explain how images and sounds attached themselves to waves or were carried by them. As an adult, I met a woman electronics engineer who told me that much of electronics was discovered by accident and was able to be used for all kinds of purpose without really being understood how it worked. That would have been nice to know when I was young.

Similarly, one time when I was young one of my uncles brought me a toy metal car about the size of a match box. He waved his hand over it and it moved. He asked me what moved it. I said he had just pushed it with his hand. He did it again, slower, showing me that he never touched it, and yet it moved. He asked me again what made it move. I had no idea. He showed me a plastic ring he was wearing that had something in it, a magnet, and said that magnets attracted metal, so the magnet pulled the car toward his hand. That didn't help me, because the obvious question for me was "How does a magnet do that?" He didn't know but thought it didn't matter. Well, it mattered to me, because the question was what made it move and how, and just telling me that a magnet made it move didn't explain the "how." I found out years later, it mattered to Einstein and other physicists too. Moreover, I still have questions about waves and physics, but now in terms of quantum mechanics, and I think it would have been both cool and helpful if I could have had an AI tutor, given that AI seems to explain physics concepts fairly well and can point out when things are not known or what the evidence is for and against any particular explanation. Given that a child could keep asking AI questions or rephrasing his or her previous ones until getting an answer that makes sense to him or her or finding out that no one knows the answer yet and the question is a serious and legitimate one is a gift to be desired.

However, none of this is totally new. Books and writing were the first instruments for remote learning. Writing and books let people communicate with each other across vast distances of space and time, although communications across vast periods of time

only went from people in the past to those in the future, and not in the reverse direction. But books do not give you the capability to question the author unless you can communicate with the author through another means, such as email or surface mail correspondence or direct chat via phone or internet. In ninth grade I argued with my English teacher—whom I adored—about her answer to a question in one exercise in the book. Neither of us could convince the other. I wrote the textbook’s author, John Warriner, and he wrote back that my teacher was correct but that that particular question should not have been in the book because the grammar concept that provided the correct answer to it was not covered until the 10th-grade book. He reported that he was going to change that question in the ninth-grade book in the next edition. He did change it; in my younger sister’s ninth-grade English course at a different high school that had a mixture of the old edition and the new edition, that one question in that exercise was different for the students who had the different editions. They didn’t know why that was and puzzled over it until my sister took in copies of my letter and Warriner’s response. “Remote” learning has always been a great thing. The internet speeds it up by making generally fruitful searches possible about specific questions; well-developed AI on the internet can make it even more efficient, a point to which I now turn.

Main AI Benefit: Always Immediately Available for Follow-up Questions

Not all books suitably explain material to all people; some books are worse than others. I and most of my classmates found college organic chemistry extremely difficult, and at the end of the term, the professor announced that the department had used a new textbook that term that arranged all the material significantly differently from the previous textbooks, and that the new book had turned out to be a disaster for this class and the department apologized. That experience helped me learn art history, though, because when I couldn’t follow the way the course textbook taught it, I found a different book that explained it all in a way I could follow. The course textbook covered painting, sculpture, and architecture in something like 50-year segments for each of three parts of Europe, and that made it so disconnected for me that I couldn’t follow it or see any threads. A second book I found organized it differently, taking the reader through painting from the Renaissance to the present in each section of Europe individually, then doing the same with sculpture and then with architecture. I could follow it longitudinally like that and then fairly easily see for myself the kinds of cross-sections timewise that the course textbook described. A really good knowledgeable tutor eliminates this sort of problem by being able to tailor instruction to the questions, needs, and understanding of the students—if the students, like Khan’s students, engage and really want to learn the material, and keep asking questions or explaining problems they are having understanding until they are finally satisfied with the information provided, at least at that time.

Khan believes AI tutors can also motivate students, which is probably true now of students he has who are curious or otherwise already motivated to some extent, but not true of many of the typical students today really not interested in academic subjects. The best students as a group that I ever taught philosophy to were Black students with low GPAs from urban areas with some of the lowest performing schools (that often did not even have working heat in the winter) in a community college in Alabama (a state that ranks extremely low in most measures of educational

achievement), particularly in Black urban areas. These students engaged in philosophy and argued better as a group than all the other groups I taught. They said they did better in my course than in previous courses because it was the first course any teacher ever cared what they thought and respected them enough to argue seriously with them till agreement, or at least understanding, was reached. Students I teach now in Alabama are nowhere near as good, but it is not because they grew up in Alabama. There is something larger involved and it affects students even from states with overall better educational systems, as will be explained shortly.

Confusing Reasonable Judgment and Disagreement with Bias

Where Khan and others believe AI needs to be unbiased in what it teaches, I fear that confuses bias with judgment. Engaged students often challenge teacher judgment and learn much of value in the process, sometimes immediately, sometimes much later. I have had many insights from reliving and resolving disagreements with teachers, one even 50 years after the initial disagreement.

But I admit that disagreeing with students these days does not seem to motivate them to think the way it used to, at least not for me, on the same scale anymore. Today's college students, as a group, though with some few exceptions, seem unable to get past whatever conventional wisdom their particular culture or subculture harbors and who seem unable to follow even the simplest of logic and do not at all like being challenged about any view they have, and consider it disrespectful or not worth having to endure. They seem to consider all judgment to be mere bias. But the earlier experience shows the problem is not because the students are from Alabama. Possibly a really brilliant and knowledgeable AI can reach and motivate such students, so I cannot rule it out, in keeping with the idea that we are talking about really well-developed (future) AI models. But I have my suspicions, unless culture itself changes, as it sometimes does and far more people want to understand things again more deeply than they now seek. And, incidentally, although not all judgments or disagreements involve what might be labeled "conservative" or "liberal" sides, I have found that on all kinds of issues, not just ones that tend to be labeled as liberal or conservative, liberal students in blue states are just as resistant to reasoning as conservative ones from red states. I have even had students argue about and refuse to believe the solution to math word problems presented because they do have results that conflict with intuition.

Questioning Khan's Socratic Questions

But this leads me to question what seems to me to be Khan's notion of the Socratic method, which seems to be the common view that just asking students any questions about the material or what they think about the material or what they think it means is the Socratic method. No, it isn't. Socratic questions are those which lead or are intended to lead students logically—not just psychologically, but logically—along a path that gets them to understand at least the problems with mistaken beliefs, and, if possible, get them to see logically what is better or more likely correct. For an example and links to further explanation about the method, see my ["The Socratic Method: Teaching by Asking Instead of by Telling"](#). Unfortunately, today the Socratic method gets a bad reputation as being a way to badger and ridicule students or is mistakenly seen to be just asking them any sort of question, such as "what do you think this story

means?” or “how would you solve this problem?” and the like. In Khan’s opening story in the book—the story about his then 11-year-old daughter Diya, collaborating with AI to write a short story—he seems to think the question the AI asked “What should happen next?” (p. xii) was Socratic. More explicitly, he considers “Why do you think the Founders included the Second Amendment to begin with?” and then later “Why do you think the Second Amendment has persisted so long?” to be Socratic questions to help answer “Why do we have the Second Amendment? It seems crazy!” (p. 22) But that is just basically repeating the question, not Socratically addressing it through anything logically leading to a reasonable answer. It is more like asking “Why do you think we have the Second Amendment if it seems so crazy?” Khan, however, uses this as an example of what he considers a challenging question with no, or much less, bias shown by the teacher. But since there are better and worse reasons on either side of the gun control issue, as with many issues, I don’t see why teachers should not challenge students’ reasoning on either side when their reasoning is flawed. Arguing against the validity of specific reasons for a point of view does not mean one is arguing for the opposite point of view. One can show that particular evidence for a claim or conclusion is insufficient without that meaning one believes the claim or conclusion itself is false. Many students understand that, and all students should.

The Socratic method when used correctly is a potentially powerful teaching method for things involving logic, but not for teaching just facts. I say “potentially” because not all people like to have to answer Socratic questions and go through the deductive process instead of just being given the answer, even when they don’t feel they are being ridiculed, though many do get enthusiastically involved with good, challenging Socratic questions, particularly when the questions cause cognitive dissonance the person feels the need to resolve. I believe Khan’s basic point is correct that if and when AI can do Socratic questioning well, it will be a great boon educationally for students around the world who want to understand things, who have access to the internet, and who cannot find through a regular internet search helpful material already prepared.

And clearly, the internet has been a major source of teaching and learning. You can find out how to diagnose and repair all kinds of things, for example, that will save you the cost of having to hire someone else to do it. The economic benefits of the internet involve not just the money it has helped people earn, but the money and time it has saved people to get information even if it meant not having to hire people who then miss out on earning the money they would have made. All of which leads to the apparently ubiquitous issue in books about AI, this one not excepted, that AI, when developed sufficiently, might be problematic by costing people jobs it does instead of them. The answers I have seen so far, including in this book, are that with regard to jobs that require insight, creativity, understanding, and the like, educators can harness the power of AI to help them do their jobs, particularly by taking over the more boring, repetitive kinds of tasks people don’t like to do and that rob them of time and energy for doing the more meaningful parts of their jobs. And they can also use AI to help them be even more creative and better teachers, counselors, or academically helpful parents. Basically, the preaching is that AI should be considered to be, and utilized as, a convenience and a collaborator, not a competitor.

Distributing Leisure Fairly

I accept all that, but I still think 1) it sells possible future AI short, if AI can indeed be perfected to think and act at least as well as humans, and 2) it misses the fact that having machines do your work for you would be (and is generally) a blessing, not a curse—as long as you earn the income from using the machine (as when you own the machine) instead of that income going to someone else (who owns the machine), or as long as your standard of living is not determined solely by your job, when there are no need for humans to do the jobs. This is an ethics economics problem about the fairest and most reasonable distribution of leisure and whatever work remains needing to be done by humans, not just a monetary economics problem nor a technology problem. Incorporating leisure into a dynamic economy is difficult for a number of reasons. (Garlikov, n.d.) Clearly if everyone's work were done by machines and AI, we could all have more leisure, which would be a good thing. And even if AI becomes sentient, as I believe it could fairly easily be made to be, at least as sentient as we are and in the ways we likely are including having as much of a sense of “self” as we do, it could still be fair to have AI do most of our work. AI doesn't tire or have limited lifespans, since its memory and other attributes can be transferred to another computer, keeping it “alive” indefinitely as the person it has been.

And also clearly, there is already so much helpful information on the internet now that is freely available in a worldwide “neighborly” kind of way, that much of the internet fulfills the Marxian goal “From each according to his ability, to each according to his needs” (Engels, 1891/2021). This Marxian goal is approached when people freely contribute their considerable skills and talents to society (in this case, their knowledge on the internet) and receive resources (the knowledge freely available on the internet) based on their needs, not earnings from their labor. Not everyone on the internet is seeking to be a wealthy influencer or seeking to monetize their contributions at all. Good AI would make even more information available and also make it much easier to access just by asking it to teach you things and explain things to you that you didn't understand, and to do it tirelessly as you have additional or rephrased questions and needs.

Can Computers Make People Be Able to Think?

Notice that the questions about AI have gone from “Can we get computers to be able to think, know, and understand things?” to “Can computers get students to think, know, and understand things?” Khan and I both answer “yes” to the second question. I probably go further than he would in answering “yes” to the first question because I think we can make machines that are sentient and that can think in the same way and at least as well as we do and have as much a sense of self as people do, whatever that is. I believe it feasible to create a computer that would have emotions, including desires and satisfactions, based on the monitoring of its internal electronic states in a way that at least seems to parallel our own perceptions about our emotional and other psychological biological (nervous system) states. And I believe that such a computer can have moral and aesthetic principles. I think there could be machines, for example, that have happiness in Aristotle's (1962) understanding of it: an activity of the soul in conformity with excellence (p. 17); and I would add, and/or the pursuit of excellence, particularly when the result is expected to be successful. And furthermore, such a computer that could experience happiness of the sort involving achievement of a

sought goal, perceived to be “desired” just for being sought, as Jane Austen (1814) described in *Mansfield Park*: “It was long before Fanny could recover from the agitating happiness of such an hour as was formed by the last thirty minutes of expectation and the first of fruition.”

References

- Aristotle. (1962). *Nicomachean ethics*. The Bobbs-Merrill Company.
- Austen, J. (1814). *Mansfield Park*. Project Gutenberg. Retrieved April 3, 2026, from <https://www.gutenberg.org/files/141/141-h/141-h.htm>
- Engels, F. (1891/2021). Foreword. In K. Marx, *Critique of the Gotha Program* (Reprint, 1st ed., pp. 1–2). Foreign Languages Press. Retrieved March 28, 2026, from <https://foreignlanguages.press/wp-content/uploads/2025/06/C16-Critique-of-the-Gotha-Program-2nd-Printing.pdf>
- Garlikov, R. (n.d.). *Ethical and philosophical foundations of economics*. Retrieved March 28, 2026, from <http://www.garlikov.com/EPFE.html#chapter27>
- Hofstadter, D. R. (1986). *Metamagical themas*. Random House Publishing Group.
- Mollick, E. (2024). *Co-Intelligence: Living and working with AI*. Penguin Publishing Group.
- Turing, A. M. (1950, October). Computing machinery and intelligence. *Mind*, *LIX* (236), 433-460. <https://doi.org/10.1093/mind/lix.236.433>.

About the Reviewer

Rick Garlikov is a philosopher and photographer in Irondale, Alabama. He holds a master's degree in philosophy from the University of Michigan (1969) and has taught for many years in colleges and universities including the University of Michigan and the University of Alabama at Birmingham. His remarkable corpus of writings on education and philosophy are freely available at www.garlikov.com.



About the Author

Salman Khan attended the Massachusetts Institute of Technology graduating in 1998 with a Master of Engineering degree in electrical engineering & computer science. He is the founder and CEO of Khan Academy, a nonprofit educational organization that offers free lessons in math, science, and humanities, as well as tools for parents, teachers, and districts to track student progress. The organization partners with more than 500 public school districts and schools across the United States. Worldwide, Khan Academy has more than 150 million registered users in 190 countries, with free lessons available in more than 50 languages.



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/> *Education Review* is free-to-read and free-to-publish.



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*.