

Brave New Words

How AI Will
Re|volutionize
Education (and
Why That's a
Good Thing) ☀

Salman Khan

'A timely masterclass for anyone
interested in the future of learning'
BILL GATES



PENGUIN BOOKS

BRAVE NEW WORDS

'It's often said that we yearn for change yet cling to the familiar. Nowhere has this been more true nor more tragic than in the field of education. Now is the time to embrace change, and to do so bravely, as Salman Khan argues in this timely treatise on AI. Leaps forward in generative AI can and should change the way all of us learn. And

I can think of no better guide to show us the way than Salman Khan' Angela Duckworth, author of *Grit: The Power and Passion of Perseverance*

'Salman Khan provides an invaluable guide for those seeking to understand and shape AI for good. This is a book about AI, education and work, but the ideas are transferable to every domain' Sam Altman, CEO of OpenAI

'Salman Khan once thought he'd become an AI researcher. Instead, he became one of the most important educators of our time. Here, with incredible intelligence and perception, he shows how learners everywhere can harness AI and other technologies for their own productivity and entrepreneurship. There are no braver words than these' Satya Nadella, executive chairman and CEO of Microsoft

'There's no doubt in my mind that Salman Khan is the leading thinker about education in the age of AI. In *Brave New Words*, he offers deep insights into how AI can be a catalyst for creativity and a tool for unleashing student potential. This book is an invaluable resource for anyone interested in the future of learning and technology' Francis Ford Coppola

'Education is the route to professional and personal fulfilment, a liberation of the mind from ignorance and a pathway to achievement. AI has today opened this possibility in new and profound ways. And rather than fear this change, by embracing it everyone can have a world-class tutor providing learning built around the needs of every individual. This requires a bold leap of imagination in what is possible in education. But as Salman Khan brilliantly lays out, this is a necessary and healthy change to a status quo that is too stuck in the past. It is a leap we must take' Tony Blair

'Salman Khan has long been on the cutting edge of education, and in *Brave New Words*, he shows just what a masterclass this book is a timely masterclass for anyone interested in the future of learning in the AI era. No one has thought more about these issues – or has more interesting things to say about them – than Sal' Bill Gates

‘A gripping guide to the future of learning. Instead of living in fear of AI, Salman Khan lays out a clear-eyed, compelling vision to empower teachers, enliven classrooms, and unlock potential in students. This book is required reading for everyone who cares about education’ Adam Grant

‘Salman Khan’s *Brave New Words* is a must-read exploration of how AI will transform education and how it can be harnessed as a force for good. The book is a timely, thought-provoking, and ultimately hopeful examination of how this powerful new technology can equalize access to learning and develop personalized instruction. We are lucky to have Khan as our guide into the exciting future of education’ Laurene Powell Jobs, founder and president of Emerson Collective

‘In *Brave New Words*, Salman Khan challenges us to reimagine the future of education in an AI-driven era. His book is not only insightful but also inspiring, demonstrating the enormous potential of technology to empower both learners and educators. Khan’s expertise and passion make this a compelling read for anyone invested in the future of education’ Arne Duncan, former United States Secretary of Education

‘Salman Khan’s *Brave New Words* is a gift to educators worldwide. This is a chance to learn from one of the world’s most visionary educators and technology entrepreneurs about how AI enables us to drive educational equity and transformation’ Wendy Kopp, CEO and cofounder of Teach for All

‘Khan provides a wealth of advice on connecting AI to education, and he writes with clarity, expertise, and optimism’ *Kirkus Reviews*

ABOUT THE AUTHOR

Salman Khan is the founder and CEO of Khan Academy, a nonprofit with the mission of providing a free, world-class education for anyone, anywhere. His lessons on YouTube have revolutionized learning, garnering billions of views globally. Recognized as one of *Time*’s 100 Most Influential People in the World, Khan is also the founder of schoolhouse.world, Khan Lab School, and Khan World School, and the author of *The One World Schoolhouse*. He holds three degrees from MIT and an MBA from Harvard Business School.

SALMAN KHAN

Brave New Words

*How AI Will Revolutionize Education
(and Why That's a Good Thing)*



Copyrighted Material
PENGUIN BOOKS

PENGUIN BOOKS

UK | USA | Canada | Ireland | Australia
India | New Zealand | South Africa

Penguin Books is part of the Penguin Random House group of companies
whose addresses can be found at global.penguinrandomhouse.com.

Penguin Random House UK
One Embassy Gardens, 8 Viaduct Gardens, London SW11 7BW
penguin.co.uk



Penguin
Random House
UK

First published in the United States of America by Viking,
an imprint of Penguin Random House LLC 2024

First published in Great Britain by Allen Lane 2024

Published in paperback by Penguin Books 2025
001

Copyright © Salman Khan, 2025

Penguin Random House values and supports copyright.

Copyright fuels creativity, encourages diverse voices, promotes freedom
of expression and supports a vibrant culture. Thank you for purchasing
an authorized edition of this book and for respecting intellectual property
laws by not reproducing, scanning or distributing any part of it by any
means without permission. You are supporting authors and enabling
Penguin Random House to continue to publish books for everyone.

No part of this book may be used or reproduced in any manner for the
purpose of training artificial intelligence technologies or systems. In accordance
with Article 4(3) of the DSM Directive 2019/790, Penguin Random House
expressly reserves this work from the text and data mining exception.

The moral right of the author has been asserted

Typeset by Jouve (UK), Milton Keynes

Printed and bound in Great Britain by Clays Ltd, Elcograf S.p.A.

The authorized representative in the EEA is Penguin Random House Ireland,
Morrison Chambers, 32 Nassau Street, Dublin D02 YH68

A CIP catalogue record for this book is available from the British Library

ISBN: 978-1-802-06489-6

Penguin Random House is committed to a sustainable future
for our business, our readers and our planet. This book is made from
Forest Stewardship Council® certified paper.



Copyrighted material

For Umaina, Imran, Diya, Azad, and Polly

Copyrighted Material

Copyrighted Material

CONTENTS

Introduction: Let's Write a New Story Together *xi*

Part I | RISE OF THE AI TUTOR 1

Throwing Away the Bottle 3

How to Teach Everything to Everyone 8

Rise of the AI Tutor 12

Part II | GIVING VOICE TO THE SOCIAL SCIENCES 25

Why Students Write 27

The Future of Reading Comprehension,
Where Literature Comes Alive! 36

AI and Creativity 41

Conversing with History 52

Copyrighted Material

Part III | **EMPOWERING THE NEXT INNOVATORS** 65

Using Science to Study Science 67

1 + 1 = Closing the Math Gap 78

Accessing Courses That Students
Otherwise Would Not 85

The Most Important Subject-Matter
Domain to Master 90

Part IV | **BETTER TOGETHER** 93

Bolstering Collaborative Learning 95

AI Meets Student Mental Health Coaching 100

The Place for Parents in AI-Based Education 108

Increasing Points of Connection Between
Parents and Their Kids 115

Part V | **KEEPING KIDS SAFE** 119

Delivering the Facts: The State of Bias
and Misinformation 121

What About Data Collection? 128

AI and the Gift of Transparency 131

AI as “Guardian Angel” 135

Part VI | **TEACHING IN THE AGE OF AI** 141

How AI Will Supercharge Teachers
and Teaching 143

Dawn of the AI Teaching Assistant 148

Copyrighted Material

Helping Build Alternative Education Models 155

Fixing Cheating in College 159

Part VII | **THE GLOBAL CLASSROOM** 165

The Global Classroom 167

Economics of AI in Education 172

Part VIII | **AI, ASSESSMENTS,
AND ADMISSIONS** 177

The Future of K–12 Assessments 179

The AI of College Admissions 186

Part IX | **WORK AND WHAT COMES NEXT** 195

Employment in an AI World 197

How to Prepare Kids to Thrive in the
AI-Future Workplace 201

Matchmaking Between Job Seekers
and Employers 207

Where This Leaves Us and Where It Will Take Us:
A Call for Educated Bravery 212

Acknowledgments 223

Notes 227

Index 229

Copyrighted Material

Copyrighted Material

INTRODUCTION:

Let's Write a New Story Together

You can't cross the sea merely by standing and staring at the water.

—RABINDRANATH TAGORE

But, whatever you do, do not let the past be a straitjacket!

—ROBERT A. HEINLEIN, *The Moon Is a Harsh Mistress*

It was New Year's Day 2023 when my eleven-year-old daughter, Diya, and I decided to spend the afternoon writing a short story together. As I set a laptop on the kitchen counter for us, Diya landed on our story's premise and then began to type. Our main character was a sassy influencer. She called her Samantha. Samantha had gotten stranded on a deserted island. "It's wonderful," I said as Diya tinkered with the story.

At which point, as if stepping into the pages of a science fiction novel or some alternate reality, Diya typed directly to her fictional character.

Samantha then responded directly to Diya, her words appearing on the screen. "Hi Diya and Sal, this is Samantha. This storytelling adventure you are writing is quite splendid! You've created a modern twist on the classic deserted-island tale!"

Copyrighted Material

INTRODUCTION

On the page, Diya typed, “How is life on the deserted island?”

“Life on this deserted island has been quite the unexpected adventure,” Samantha replied. “I’ve had to learn how to survive without my usual comforts, but I’ve also discovered some hidden talents I never knew I had!”

“What kind of talents?” Diya asked.

Samantha replied, “For one, I’ve become quite skilled at building makeshift shelters using palm leaves and branches. I’ve also learned how to catch fish with my bare hands and start a fire using only sticks and dry leaves!”

Without missing a beat, Diya wrote, “Why don’t you use your scarf to make a hammock?”

“Use my scarf? No way!” Samantha replied. “This is a designer piece, and I can’t risk getting it ruined. Now, what do you say we keep writing this story I’m in together? What should happen next?”

“I think we should do what Samantha says,” I said.

Diya agreed. Returning to the computer, she typed, “What if, all of a sudden, Samantha’s best friend, Emily, appears on her sailboat. She’s here to rescue her.”

“I love it!” Samantha wrote the next part on her own. “Samantha couldn’t believe her eyes when she saw her best friend, Emily, sailing toward the island on a small sailboat. She waved her arms frantically, trying to catch Emily’s attention. As the sailboat approached the shore, Samantha’s excitement grew. She couldn’t wait to tell Emily about her adventures on the island and hear how Emily had managed to find her. Okay, now it’s your turn again, Diya and Sal! Write the next two sentences, and let’s see how their reunion unfolds.”

So, to recap: my eleven-year-old daughter and I were now having

Copyrighted Material

a real conversation with a make-believe character. This was some Kurt Vonnegut-level absurdity. The world as I knew it had suddenly shifted toward a strange new reality.

It would never be the same again—not for us, and not for anyone else on the planet.

The Offer

To explain how we got to this incredible scenario, let's back up a bit.

Twenty years earlier, 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 mastering. I called the website the 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,

Copyrighted Material

INTRODUCTION

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.

My long-held aspiration for the organization was that it would act as a tutor for every learner in the world, an endeavor that has always been our true north. This wasn't simply a matter of scaling personalized support for the sake of it. Long before Khan Academy, decades of research (and intuition) pointed to the idea that kids could learn much more if the pacing adapted to the student and allowed each one to truly ace a subject (that is, mastery learning). This contrasts with the status quo, where a class of thirty students often moves on to the next concept even when a good portion of the students haven't yet displayed proficiency. Obviously, getting every student a dedicated on-call human tutor is cost prohibitive. The only viable solution is to work with technology. It seemed to me that AI technology might someday turn out to be an important part of that puzzle, maybe even the holy grail of truly emulating a real tutor.

I am not alone in having this dream. The science fiction writer Neal Stephenson wrote about the potential influence of technology on education in his novel *The Diamond Age*. The book is set in a future world that introduces the concept of using AI, in the form of a highly advanced interactive book and app called *A Young Lady's Illustrated Primer*, to provide personalized education to its young

users. The Orson Scott Card novel *Ender's Game* imagines a battle school that employs advanced AI technology to test and train students' strategic thinking and decision-making skills through a personal AI tutor called Jane. Isaac Asimov's short story "The Fun They Had" describes a school of the future that uses advanced technology to revolutionize the educational experience, enhancing individualized learning and providing students with personalized instruction and robot teachers. Such science fiction has gone on to inspire very real innovation. In a 1984 *Newsweek* interview, Apple's co-founder Steve Jobs predicted computers were going to be a bicycle for our minds, extending our capabilities, knowledge, and creativity, much the way a ten-speed amplifies our physical abilities. For decades, we have been fascinated by the idea that we can use computers to help educate people.

What connects these science fiction narratives is that they all imagined computers might eventually emulate what we view as intelligence. Real-life researchers have been working for more than sixty years to make this AI vision a reality. In 1962, the checkers master Robert Nealey played the game against an IBM 7094 computer, and the computer beat him. A few years prior, in 1957, the psychologist Frank Rosenblatt created Perceptron, the first artificial neural network, a computer simulation of a collection of neurons and synapses trained to perform certain tasks. In the decades following such innovations in early AI, we had the computation power to tackle systems only as complex as the brain of an earthworm or insect. We also had limited techniques and data to train these networks.

The technology has come a long way in the ensuing decades, driving some of the most common products and apps today, from the

INTRODUCTION

recommendation engines on movie streaming services to voice-controlled personal assistants such as Siri and Alexa. AI has gotten so good at mimicking human behavior that oftentimes we cannot distinguish between human and machine responses. Meanwhile, not only has the computation power developed enough to tackle systems approaching the complexity of the human brain, but there have been significant breakthroughs in structuring and training these neural networks. One of the more recent breakthroughs came in 2017 with the advent of transformer technology from Google, which allows for, among other things, better and faster training and more accuracy in how words and ideas connect based on this training information.

How good these systems can get is usually related to the complexity and architecture of the underlying “model.” Think of a model as a computational representation that tries to mimic or simulate something in the real world. For example, when meteorologists try to predict the path of a hurricane, they use weather models that contain a software representation of billions or trillions of smaller volumes of the atmosphere and forecast how those smaller volumes would likely interact with one another. In the case of large language models, they are specifically designed to model associations between words. In this case we’re modeling not atmospheric conditions but neurons and synapses. Large language models such as GPT-4, short for Generative Pre-trained Transformer, are essentially big, powerful—albeit digital—“word brains” trained on a colossal amount of information from books, articles, websites, and all sorts of written material.

By analyzing and processing this vast amount of text, the language model learns the patterns, the language, and the context of

how words, sentences, and paragraphs fit together. If you were to ask a large language model like GPT-4 a question, it would know what to reply based on its training from all those books, web pages, video transcripts, and social media posts. What it lacks in real-world sensory experiences of the human brain, it compensates for by having exposure to more language than any human might hope to read, watch, or listen to in multiple lifetimes.

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 groundbreaking 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. I didn't realize it yet, but the world was about to be turned upside down.

For context, OpenAI was still four months away from releasing ChatGPT and seven months away from releasing GPT-4, the eventual rollout of which is what they wanted to discuss. I was intrigued but skeptical that there would be anything we could do together. I didn't have a strong view that some of the newer-generation generative AIs would be immediately applicable to our mission. Advances in AI technology had already done some interesting things around writing that sounded credible, but in my mind the technology didn't yet seem to have a real handle on knowledge and also lacked the ability to perform logical or deductive reasoning or reliably produce legitimate facts. At the same time, however, I had a lot of respect for what OpenAI had already accomplished. So we scheduled some time together and met.

Each successive generation of these models typically had far more complexity, bluntly measured by the number of parameters

INTRODUCTION

they contained. The best way to think of a parameter is a number describing the strength of a connection between two nodes in the neural net that represent the large language model. You can view it as a representation of the strength of a synapse between two neurons in a brain. When it was first launched in 2018, GPT-1 had more than 100 million parameters. Just a year later, GPT-2 had more than 1 billion. GPT-3 had more than 175 billion. GPT-4 was likely to have on the order of 1 trillion parameters.

The OpenAI leadership felt that GPT-4 was going to surprise folks with its enhanced capabilities, which they believed would both excite and possibly unnerve a lot of people. Because of this, they were looking to launch it alongside a small number of trusted partners that might be able to showcase socially positive and real-world examples; Khan Academy was the first organization that came to mind. The second reason they wanted to reach out to us was to help evaluate the AI itself. They needed to show that GPT-4 was capable of deductive reasoning, critical thinking, and actually dealing with knowledge. The OpenAI team aspired to see how GPT-4 would perform on college-level biology questions, and we had thousands of these.

I found myself suddenly exhilarated to be one of the first people on the planet to see the capabilities of GPT-4. Based on past experience, I knew the time to really explore a technology is when it's on its way toward getting good. If you invest and test it properly when most still believe that it's a toy, or a distraction, you put yourself in position to really reap its benefits when it is ready for prime time. This was the case in the infancy of video learning, when plenty of naysayers said YouTube was simply an idle pastime. But early pioneers showed us that on-demand videos were so much more than

cats playing piano and that, in fact, you could use them to help people learn.

Today, it's commonplace for students to learn almost anything they want using on-demand videos, and it's become much more accepted in the classroom. Khan Academy has played a leading role here, using video to support on-demand help for hundreds of millions of learners around the world. We've also shown that rather than somehow being a substitute for the teacher, videos can off-load pieces of a lecture, freeing up more time for personalized learning, hands-on activities, or classroom conversation. This arguably makes the teachers more valuable, not less. And now it was time to see if generative AI could do the same—support students and let teachers move up the value chain.

Sam and Greg started their GPT-4 demo by showing me an AP biology multiple-choice question they had pulled directly from the College Board's website. They asked me for the answer. After reading through it, I said the answer is *C*. They then asked GPT-4 to answer the question using a chat interface (similar to what folks are now used to with ChatGPT). A moment later, GPT-4 answered the question correctly.

I didn't say anything right away, but secretly I started to get goose bumps, even as I remained somewhat skeptical. "Wait a second," I said. "This is an AI that can already answer a biology question at an AP level?" Maybe it just got lucky on this example, I thought. "Can you ask it to explain how it got the answer?"

Greg typed, "Please explain how you got the answer." Within seconds, GPT-4 provided us with a clear, simple, and thorough explanation. Not only that, but it was so conversational that it might as well have been a human being answering, not a machine.

Copyrighted Material

INTRODUCTION

At this point, I gave up on hiding my amazement.

“Can you ask it to explain why the other choices aren’t correct?”

Greg obliged, and a moment later GPT-4 explained why all other answer possibilities in the AP question were wrong.

Next, I asked Greg if GPT-4 was capable of writing an original AP-level question.

It did, and then it wrote ten more.

Two months later, I visited Bill Gates to give an update on Khan Academy, and I learned why the OpenAI folks had shown an AP biology question. Bill told me that when he first encountered GPT-3, he was impressed, but he had told the OpenAI team that he’d only be *really* impressed if it was capable of passing the AP biology exam. What the OpenAI team had shown me in that first demo was that GPT-4 could now do so.

“This changes everything,” I said to Greg and Sam, my mind spinning with possible ways that GPT-4 might allow us to reimagine education, credentials, work, and human potential.

“We were kind of thinking the same thing,” Sam said. “It is not perfect yet, but the technology is getting better. Who knows? If we get it right, this might be something educators will want to use.”

Technology that we had until recently thought of as something straight out of *Star Trek* was suddenly very real. The innovation imagined by the greatest science fiction writers had become a reality.

Time for a Hack-AI-Thon

In the early 1940s, Claude Shannon, a brilliant mathematician, produced several consequential theories. Among them, he mapped out a

theory of electronic communication that would become the basis of digital technology. In 1948, while working for Bell Labs, he started dabbling in the field we now know as artificial intelligence. Shannon decided to play with how an algorithm approximates language. He published a paper in *The Bell System Technical Journal* called “A Mathematical Theory of Communication.” It was the early days of digital computers—well before the advent of the internet—and Shannon’s information theory first made the case that a series of probabilistic processes could approximate the English language. By keeping track of how many times words appeared in a text, he devised an algorithm where he was able to predict what word was most likely to come next. Eventually, this small language model generated a sentence. The better this process got, the more natural the sentence sounded. It’s an oversimplification, but the likes of GPT-3 and GPT-4 are essentially far more complex large language models based on training a neural net in very specialized ways, and the underlying idea can be traced back to this early work by Shannon.

Soon after the development of Shannon’s work, we saw another great mind enter the realm that was to become artificial intelligence—a computer scientist by the name of Alan Turing. Above and beyond his work breaking German codes and helping us beat the Nazis, Turing explored the concept of AI and whether machines can reach a point where they can mimic human intelligence convincingly. In 1950, he wrote a foundational paper titled “Computing Machinery and Intelligence,” where he introduced the concept of the imitation game, which we now know as the Turing test. Imagine you are having a conversation, but you can’t see the person you are talking to. It could be a person you are chatting with on your computer or phone. Now, if you can’t see or physically interact with the other person,

INTRODUCTION

how can you tell if you are talking to a human or a machine? That's the essence of the Turing test. To perform the test, there's usually a judge involved who is responsible for evaluating the responses from both the human and the machine. The goal for the machine is to convince the judge that it is, in fact, a human. It needs to demonstrate intelligence, understanding, and the ability to hold a coherent conversation, just as a person would. Turing proposed that if a machine could consistently fool the judge into thinking it's human, we could consider it intelligent. In other words, if the machine can pass the Turing test, it would imply that it possesses humanlike intelligence.

When I accepted Sam and Greg's offer to field-test the new GPT-4 technology in the summer of 2022, I wondered how close it might be to passing the Turing test. I had studied artificial intelligence at MIT in the mid-1990s. Back then, there were simple programs that could trick a person for a few interactions but nothing that could consistently feel human across a long, detailed conversation. It always seemed fantastical that one day a machine might actually pass the Turing test, let alone in my lifetime, and it was thrilling to actually try out a technology that seemed on the cusp of passing it, or that perhaps had already done so. This advancement could be tantamount to scientists achieving cold fusion or faster-than-light travel.

As the initial wave of excitement crested, I also started thinking about the societal implications of a seemingly intelligent technology. While capable of solving so many problems, AI was also capable of introducing some potential downsides. If this large language model was going to be able to help tutor students, then it was going to be

Copyrighted Material

capable of writing essays for them too. What if the new version of GPT managed to be only a crutch for our students, preventing them from developing their own research and writing skills? It also occurred to me that if GPT-4 had the capacity to empower folks by helping them communicate and solve problems, then it also potentially threatened to dislocate many people from their jobs and their sense of purpose. A convincingly human technology capable of being a great tutor might also be a technology that bad actors could use to defraud or brainwash unsuspecting people.

My mind continued to spin a great variety of dark scenarios and outcomes—from data collection on our kids to the potentially addictive quality of the technology. I understood that the disruptive nature of the AI meant that we all needed to take it seriously. Beyond OpenAI, there are several other organizations investing heavily in large language models, including Microsoft, Google, and Meta, not to mention state actors like Russia and China. All of the technology giants had been using some form of artificial intelligence for years to feed us the ads, videos, search results, and social media posts with which we engage on an hourly basis. But this AI seemed different—because it was. Science fiction authors have always drawn the distinction between an artificial specialized intelligence that can optimize one thing versus an artificial general intelligence that can reason across many tasks like a human being, the latter case being the one that could lead to both utopian and dystopian futures.

To many, large language models like GPT-4 approach artificial *general* intelligence because they can write about, and seemingly reason about, any subject, making them relevant almost anywhere. Generative AI can also construct and make sense of images. While it

Copyrighted Material

INTRODUCTION

was clear that generative AI was going to change our world in ways we might scarcely imagine, it also occurred to me that it was not up to me or Khan Academy to decide whether we were comfortable with this technology moving forward. It was already moving forward. As someone trying his best to use technology for good, I wondered if we might be able to utilize generative AI to achieve the greatest positive effect and edge us closer to that utopian scenario, especially in the realm of education.

With unexpected access to OpenAI's newest model, I sat down at my computer on that very afternoon, a Friday. I was glued to it for most of the weekend. After ten hours planted at the computer, I tried to shut off and sleep, but my mind was racing at all that the technology was capable of and how we, as educators, might use it with our learners. I saw that the implications of GPT-4 were nothing short of revelatory. Not only was it able to solve complex problems and answer tricky questions, but it was also able to explain its reasoning and generate new questions on its own. It was able to take on personas with humanlike characteristics. After prompting GPT-4, I got it to act as a math tutor, a science tutor, and a language tutor, each one possessing the ability to drive a conversation.

What really caught my attention, though, was its ability to write in different tones and styles. For instance, as part of my GPT-4 learning journey, I could not resist the temptation to have GPT technology rewrite the Declaration of Independence in the voice of Donald Trump:

Copyrighted Material

Folks, let me tell you, it is time for us to declare our independence from this failing government and this terrible leader, King George III. I mean, this guy is a total disaster. He's been taking away our rights, he's been ignoring the will of the people, and he's been more interested in lining his own pockets than actually helping the people. It is just sad, really, really sad.

But we, the people, we're not going to stand for it any longer. We're going to make our own country, and it is going to be the best country, the greatest country the world has ever seen. Believe me, nobody knows more about making a country great than I do. We're going to have the best economy, the best military, the best everything. It is going to be tremendous.

Next, I asked ChatGPT to write the Declaration of Independence in the style of Dr. Seuss:

In the land of the free and the home of the
brave
A new nation was born, no longer a slave
To the rule of a king across the sea
We declared our independence, and shouted
with glee

With a pen in hand and a fire in our hearts
We wrote down our grievances, every last part
No more taxes without representation
No more soldiers in our homes, no more
frustration

We hold these truths to be oh so true
That all men are equal, yes me and you
We have rights that can't be taken away
Life, liberty, and happiness, we're here to stay.

Copyrighted Material

INTRODUCTION

Tricks like this became somewhat commonplace once people had access to ChatGPT, but at the time I was one of the first to have access to this type of technology and attempt such prompts. And GPT-4 was frankly much better than the first iteration of ChatGPT, which wouldn't be released to the public for another few months. I found the results of my prompts to be impressive, hilarious, and even a little scary. When I asked it questions or requested recommendations, it gave me answers that felt completely authentic. There was no person behind the scenes typing these responses, nor was there an algorithm generating the text with traditional if-then-type logic that other algorithms typically worked from. I did not get canned and robotic answers either. Instead, I received varied responses to the same query every time I asked it, responses that took into account the context of my conversation with it up to that point.

More specifically, it made me realize the potential of this technology to transform the way we think about K-12 and higher education and beyond. The AI was not quite perfect yet. It was getting math incorrect more than I liked, but I could even see improvement as I got better at prompting it. By the end of that weekend, I wondered what might happen if I gathered dozens of the brightest minds in technology and education to play around with the platform alongside me. OpenAI agreed to grant access to another thirty or so engineers, content creators, educators, and researchers on the Khan Academy team to experiment with GPT-4.

It was time for a hackathon.

Every six months, we have a week at Khan Academy where employees are allowed to work on pretty much anything they want related to our mission. I presented GPT-4 to a small subset of our

Copyrighted Material