



MACRO Voices
with hedge fund manager Erik Townsend

Matt Barrie: The awesome power and risk of Artificial Intelligence

August 24th, 2023

Erik: Joining me now is Matt Barrie, CEO of [Freelancer.com](https://www.freelancer.com). Matt, it's great to get you on the show, I want to talk about artificial intelligence. It's a topic that I've been aware of literally, since I was a teenager. I got my start in life, sneaking into the Artificial Intelligence Laboratory at MIT in the 1970s. So, needless to say, a long time ago. AI has been progressing very, very slowly. Since then, steadily slowly, then all of a sudden, about a year ago, something happened. And the progress has been absolutely mind boggling. What's going on?

Matt: Well, there's a few things going on. And certainly the last 6 to 12 months have been a wild ride. Probably start off, before we get into the underlying technology, is talking about some of the practical applications and kind of what's blown up there. You know, if you fast forward back, as recently as August 2022, last year, you're an illustrator, you're going along fine businesses as usual, you're getting commissions to do art, and a software tool and a stable diffusion gets released. Now, I, like many other people download several diffusion when it first came out. And this is a software package that will take a text string and turn it into an image or an illustration. And you download it, you play around with it for a few hours typing in various text strings and tweak a few parameters and you'd kind of get out the other side, some dream like sort of images, part Picasso, part Salvador Dali, you could kind of see that that image, somewhat looked like the string you'd put in as the input.

But it was nothing really more than the toy at that point in time. You'd play with it for three or four hours, you'd maybe get something that looks a bit artistic, and you'd go ok, that's interesting, but let's move on. Fast forward, two and a half months later, and another tool came out called Midjourney. And this time, you would type in a text string, and something magical would happen. You would get highly artistic, almost photorealistic images come out the other side of anything you type in, anything. And all of a sudden, the entire design world changed. And subsequent releases came out of this tooling and it's been absolutely incredible. You can type in, my niece can type in any sentence now and get out a hyper realistic high resolution image of anything. And this is like an atom bomb dropping on the design industry and subsequent evolutions of things like Adobe Firefly. And you can now take a graphical image of someone's face, which is in two dimensions. and you can twist it and turn it around and see the back of the head, even though the back of the head was never captured, in any photograph ever. You can take an image of a scene and you can zoom out of it. And the AI will infill in full photo realism, as you continue to zoom for as long as you like. Something crossed, what I call

the uncanny valley, where, this is a concept in computer science, where if a human sees a picture of another human, and it's very, very primitive, you've got very low affinity with that image, then over time, as the quality of that picture of a human improves or might be a robot, something in the physical world is the quality of the representation of a human improves, we build affinity, then it gets to a point where it becomes very unsettling. For example, in a horror movie, you see obviously something that's quite fake, but it looks like a human but not quite a human. And then you have like this big trough where the human brain feels very much at ease when it sees a picture of that human because it's kind of close to a human, but it's not actually a human. And then when you come out the other side of that uncanny valley, where you look at an image and it's photorealistic. And you know, it's human, it's real, and you build that affinity back up again. And something has happened in design where the AI was producing images that kind of look okay, not great, what have you and then all of a sudden leapt forward and we're hyper realistic. And illustrators are, they're scratching their heads going, do I have a job still, because the software now can illustrate as good as any illustrator out there can take a photo is, you know, representation of a photo as good as any photographer out there in the world. And there are serious questions about what the future is of the design industry. And this sort of leap that's happening in artificial intelligence is happening not just in images, it's happening in all sorts of other modes of communication. For example, text and chat. I'm sure many of your listeners have already played with chatGPT, right? ChatGPT is a chatbot where you type in a sentence and it will chat to you and you can ask him to write an essay for you. You can ask it to solve the problem for you, and something happened.

Also around the same time, where it went from a toy where you would ask him a question, you know, like write me a poem about a particular topic, and it wouldn't rhyme or it wouldn't really make too much sense because it was trying an interesting little novel toy, to something now where you can get chatGPT to write at a level that exceeds probably the average human by substantial amount. It can write for you a stock purchase agreement for a series a stock subscription, or financing that's happening, you can get it to answer questions that are at the level of university students would face in their assessments. It's now passing the top 1% and verb on the graduate recruitment exam, which is to get into college in the US. It's in the top 7% of the LSAT, it can pass the bar exam to become a practicing lawyer. And it's doing this in leaps and bounds across all ranges of human endeavor. And it's doing it in a way which is really magical and incomprehensible, unpredictable to the inventors that we've been working on the AI and frankly, quite shocking,

Erik: Matt, what happened that caused the pace of progress to change so quickly? We went from progress on AI every year, steadily enough to all of a sudden, leaps and bounds. Actually, I think one of the imaging companies that you talked about, went through from version 1 to version 2 to version 3 to version that ends up at version 5.1. Version 5.1 came out like four months after version 1. So the software development cycle has been shortened up is that have to do with the AI programming itself?

Matt: That's right, there's been a fundamental breakthrough in the underlying artificial intelligence technology. I mean, that particular company you're talking about is Midjourney,

Midjourney version 1, you'd say, draw a picture of Erik Townsend and you get kind of a stick figure. Version 2, you'd get a nice little coloring in that looked like maybe a five year old did it. Version 3, you would probably get the art you'd see from an average artist down at like a fair. Suddenly, the leap from version 3 to version 4, and you've got hyper realistic photo quality picture of Erik Townsend. And then version 5 came out and was even enhanced from them. And now it's going in all sorts of crazy directions, you can open Erik's mouth, you can lift his hand up, you can twist him around. And it fills in all the detail. The key innovation that happened here, effectively in the space was a breakthrough by Google called the transformer. And what the transformer did was, it allowed artificial intelligence models to consume large amounts of data to train. And as it consumed more of that data, it got substantially better.

So what essentially is happening is these models take a large data set. And by large data set, I mean something like 10% of the entire Internet. And after training on that data set, you can give it a new input sequence, say an input sequence of text, and effectively what they do is predict the next word. So I can fit it in a sentence and predict the next word, or I can ask it a question and, and it will predict the next word of an answer. And then the next word after that, the next word after that, the next word after that, and so forth. But something really magical and incomprehensible is all of a sudden happens, at least to humans. Because as you increase the model scale, and by that I mean the size of the training data, the number of parameters you've put in the model, and so forth. It's suddenly leaping forward in a level of quality and sophistication of the output. That is not predicted by the human mind, in terms of thinking in a linear progress. So you see an image it looks like a stick figure, you see an image, it looks like a primitive painting, you see an image it looks kind of okay, and all of a sudden, it just leaps forward to photorealism. The same thing with chatGPT, you know, the earlier models of chat GPT, you'd ask questions, it was like you'd expect talking to a computer program, you get a response, it was not great. It didn't really reply, as you'd expect a human to, and then all of a sudden, it leapt forward to being able to pass the GRE in the top 1% of verbal, it leapt forward all of a sudden to be able to write any document on any topic, any essay on any topic, perfectly.

Erik: Listeners, Matt has put together a write up describing his own research and experimentation with AI tools. We've got the download link in your research roundup email, I highly recommend that you download it. Matt, one of the things that I found most fascinating when I read this, is that you describe, I don't remember which one of these products it was, but it's something along the lines of the drawing the pictures and you go from the stick figure to the okay picture to the pretty good picture and eventually you get to the really good picture. Okay, it's pretty easy to understand. The goal of this thing is it's drawing pictures and over time it gets better at drawing pictures, that all make sense. What I was fascinated by is, if you're giving it the instructions, you're telling it how to draw pictures in English, it only knows how to speak English. But at some point, you get to version 5 or something, where it's drawing really good pictures. And all of a sudden, it speaks fluent Arabic. And you can tell it to draw pictures in other languages and it knows what you're talking about, what the heck is going on?

Matt: Well, the actual researchers who are developing this technology don't know, they didn't expect these emergent abilities, they didn't predict them. And they certainly, at this point, are

throwing their hands up in the air, in terms of describing how they're emerging. I mean, what's happening is that, as you increase the model scale, so you increase the amount of training data, you feed in the number of parameters. So for example, chatGPT, or the underlying model GPT went from 1 billion parameters to 175 billion parameters to 1 trillion parameters. So there's big step ups in terms of order of magnitude. But as those steps up in all orders of magnitude happen in terms of the model, the model starts to do some things that you don't expect. So for example, you give it this huge training corpus of English language texts off the internet. Of course, in the English language text, there's going to be other languages in there, because you're not actually going through and curating it heavily. You just basically got a giant web scraper that scraping down the internet, but you're training it to answer questions in English, that's what the training is doing. So you're training it, and then you'll increase in the model scale by an order of magnitude of training data, an order of magnitude in terms of complexity of the model. And you step it up, you step it up, you step it up. And just all of a sudden, this model that you've trained in English, starts answering questions in Persian perfectly. And nobody knows why. And we're seeing all sorts of emergent abilities, whether it's manifested in the image, that is drawn in Midjourney, doesn't look that realistic, and then boom, suddenly, it is photorealistic. You're seeing these, just these leaps that are happening in the last 6 to 12 months, where, as the models have managed to scale up in terms of compute power, in terms of complexity, these magical things are happening. The model doesn't know arithmetic, chatGPT can't do math, all of a sudden solving University math level problems. And possibly, what is going on there as an explanation is, similar to what happens in the human brain when the human brain learns things, right? That, some sort of symbolic system is building in terms of these models.

So effectively, all these models are, they're giant matrix multiplications, I think one of the inventors of RLHF, which is Reinforcement Learning for Human Feedback, which is basically, the AI outputs a few answers, and a human will come along and fine tune it by saying, that answer is better than that answer, etc. and so forth. And you've got these big farms of, of effectively, freelancers, now that are looking at the output of these models saying, left is better than right, left is better than right, right is better than left, etc. and so forth, in that surprising effect. But the inventor of that particular technology used to try and do the fine tuning of the models, such as OpenAI models, said that, you know, effectively what's happening in these AIs is you've got giant matrices of numbers, you're applying a non-linearity, you're then multiplying by another giant matrix of numbers. And you're doing that over and over again, maybe 150 times and you get an output. But you don't really know what's happening in each of those matrix multiplications. You can make a guess of kind of what's happening, you actually don't know what's happening when it's training, and all this data, you don't know what's going on under the hood. But what might be going on under the hood is kind of what happens in the human brain when the human brain learns something. So for example, when you learn to drive a car, as a human, you've got a bunch of senses that you're using to kind of train yourself to drive the car. You've got an instructor maybe talking to you in audio, you've got what your eyes see in terms of the images are there in front of you as of the road, and then there's maybe a car approaching and there's an obstacle on the left or a sign saying slow down or speed up or turn left, etc. So you've got the visual input, you've got the audio input, you've got the tactile input. If you're holding the wheel and your foot in the accelerator and so forth. You know, when you're

learning to drive a car, at the very early stages, it's a very conscious decision. You're very much thinking and overcorrecting and over, and you can train yourself, you know, if I jerked the wheel too much to the left, the car moves a bit too much left, then I've got to turn it back to the right again, et cetera, and so forth.

But as you get enough training data, and the instructor is talking to you for a while, you're going to experience on the road, you're getting your hours up. What actually happens at some point is when you think about turning left, you just end up turning left, you're not thinking actively. I've got to move my hand two centimeters down, left and spin something or other and put my foot on the brake or the pedal or what have you, it just kind of happens. And maybe what's kind of happening in these models is, they're just so big in terms of the number of parameters. And they've seen so much data that some were hidden away in those matrix multiplications abilities of forming it the model somehow. And the training is figuring out what turning left means, what turning right means, the very nature and the very structure of the world around it, you know?

And it turns out, at least the theory right now, with some of the inventors of this technology, such as Ilya Sutskever, who is the chief scientist of open AI is, it turns out that in order to predict the next word in a sentence reliably well and convincingly well, you actually have to know a lot about the actual world itself, more than just pausing text and pausing English. You have to kind of know, you know, what people are thinking in a room, how they feel about certain things, the physics of the world in the room, etc. and so forth, right? Otherwise, you can't get a convincingly good output. If I ask it a question, for example, how do I balance a laptop and nine eggs, a nail and a bottle, you'll get gobbledygook answer, unless you kind of know something about the underlying physics of the situation. If I want to get it to write a fiction story about two characters talking to each other, I think I sat down with you over lunch once and we talked about, let's get a fictional interview of you talking to the Russian Foreign Minister, and doing a fictional interview on MacroVoices, right? In order to have a conversation that looks convincingly real produced from the models, it's got to know about the world, it's got to know how you would think about a response that the Russian Foreign Minister gave you, you'd have to think about the biases that might be in his head in terms of how he would respond to a question that you would ask etc. and so forth. And what's being theorized here is that these models, even though they've got a stencil, they're very, very simple task, you know, predicting next word, predict next pixel to draw on a scene, etc. and so forth. What's actually happening is they get, they're developing these trillion parameter models that are developing a very, very, very good understanding of the world, that they're in the environment, the physics, the thoughts, you know, the theory of mind, or how someone feels what they think, etc. and so forth.

And as a result, because of the scale, this is kind of just emerging. These little bits and pieces are kind of emerging, the ability to do arithmetic, the ability to speak in Persian, the ability to, you know, to understand how the very nature of illustration works, not cutting and pasting images together, but the actual theory of illustration. And that is really surprising people. It's surprising, not just the consumers, they're using the tools in terms of the quality of the output. It's surprising the actual developers, the tools, who can't explain how these models are working. It's just that at a certain scale, they're crossing the uncanny valley. And you're being convinced

you're talking to a human, you've been convinced you're seeing a photograph, you've been convinced that that deep fake of Tom Cruise is actually Tom Cruise in the video, and not a fake. And that's what's happening in the last 6-12 months. And it's been doing it at an incredible pace. And there's been a few things that have kind of really driven that one is obviously the ability to script on a smaller scale. You know, large investments that kicked off open AI, from Microsoft at \$100 million, before that with Elon Musk, because these models take a huge amount of money to train, I think it's been estimated that a training run of GPT4 takes \$100 million to actually go through the training. You know, to train the Tesla car to drive takes 70,000 GPU hours of training and very, very high end array of processes, right? And so, the scale has got there. And then the underlying technology of the transformer model has allowed the AI to train at scale, because the transformer, effectively what it does is allows to look at data across great distances from other data. So in the past with neural networks, when you will kind of break into MIT lab and so forth, you had to kind of feed in data sequentially. And it had to kind of remember where it was. And as a result, it couldn't really remember that much. And it was very slow to train these models.

What happens now is, it can kind of jump around the data because it knows that the position on coding of all the data at scale, kind of when it comes into the model. And so it's allowed it to basically look across data at large distances. And combined with compute power, combined with the training sets, these emergent abilities have kind of come out of the models and it's just doing it in a way that is just mesmerizing everyone because it's not in a linear fashion. It's a super linear fashion. It's leapt forward with such an improvement that it's startling. It's startling its very creators. And you're getting some very, very interesting responses from like the godfather of effectively neural networks, Geoff Hinton who, who suddenly quits his job at Google, and proclaims, I had to leave Google to be able to talk about the dangers of AI. And by the way, I potentially regret my entire life's work.

Erik: Well, we'll definitely come back to the risks and dangers before we're done here today. But for now, I want to stay focused on this idea of the acceleration and the acquisition of these new skills. Now, if I'm understanding this correctly, the inherent nature of anything that's growing exponentially is that it tends to befuddle humans, because our brains are not good at processing exponential relationships. So what happens is, it seems that things are going very, very, very slowly, painfully slowly, then they finally start to pick up and then at the end, it's an acceleration that just you can't believe how quickly it's going. But the thing is, Matt, if this is, in fact an exponential progression, in other words, we are increasing the pace that these models are learning exponentially. What that means is that the last six months only seem fast now, and the next six months are going to bring even more faster advancements to this space, because we're seeing an exponential progression here. Is that what's happening? Or Am I misinterpreting that exponential relationship?

Matt: That could very well be what's going on. It could be that we are at that sort of knee of the exponential curve, I mean, that there's a classic example..

Erik: But doesn't that then imply that we're not very far away from the point, where humans can't even perceive or keep up with the complexity of what's going on? The AI is growing, and transforming and becoming something different than its prior self, in a way that's occurring so quickly that we humans can't even tell what's happening.

Matt: So what we're talking about here is effectively the essence of the singularity. And the singularity is a concept that Ray Kurzweil and other proponents have talked about, where technology over time accelerates. And while it can be quite slow over time, in terms of how we perceive it, I'm pretty sure that everyone, at least in the last 20, 30 years is through the advent of computer and now AI, people I think, can perceive that the world is moving faster and faster around us in terms of technology. And the thesis around singularity is at some point, the advances in technology outstrip the ability of the human mind to comprehend it. And you get this sort of limit up in terms of scientific advancement and technological discovery that moves so quickly, that humans can't keep up. And at that point, the singularity happens. And then there's this whole thesis about, at that point, does the human mind and computer somehow merge, which is what Elon Musk is trying to do with neuralink, find a bridge, where the wetware of the human brain and the software and the hardware of computer technology can come together. And so that through that singularity, we kind of merge together as one.

Or, we'll end up with something else, some other sort of scenario where, perhaps we're like the primitive fungus and humans suddenly come along, and all of a sudden, were exterminated simply because that were irrelevant in terms of the ability for us to interact in this new technological enlightenment. It could very well be that we're at the knee of that curve. I mean, you know, the human brain is not very good at predicting exponential things, right? It's designed for linear things. It's designed, over years of evolution, you see the cheetah running and the cheetah goes by the hill, your brain is designed to do the linear interpolation of the cheetah is going to come out the other side of the hill and maybe come after me, you know, I should do something, right? It's not designed to predict things that are that are exponential in nature. I mean, there's a classic example of Craig Venter when he was doing the Human Genome Project, which was sequenced the human genome, which at the time was thought to be impossible. You know, he was many, many, many, many years into that project, I think was about a decade into like a 15 year project, or 13, or 14 years into a 15 year project, where he'd only sequenced 1% of the human genome. And everyone said, this, this whole project is going to fail. But it went from 1% to 2% to 4% to 8% to 16%, 32%, 64%. And boom, it happens so quickly. And it was successfully achieved, but people didn't see it coming even well into that particular program. So it could very well be that this is what is happening with crossing the uncanny valley here, with technology, is this just that we've got to a certain amount of complexity and scale in the ability of these technologies to train and they've done so in exponential leaps.

And as a result, we're now going through that knee where, wow, it suddenly goes from what we predict the output to be to geez, I don't know what's going on anymore, and we could be heading very much into that sort of exponential explosion, I don't know. Now there are natural limits there. There are some things that might limit this ability to continue for a lot longer. One is,

how much data is there publicly on the internet? I believe GPT4 sucks down about 10% of the public data on the internet. So when it gets to 100% of that public data to the internet, you know, where does it go next? Now, obviously, there are many other data sources, they haven't been tapped into the private data sources. I mean, when you listen to the news for open AI, in addition to the web scrapers, they've also licensed a lot of private datasets. There are also a lot of datasets that haven't been tapped into yet, in other modalities, we haven't looked at video content, for example, in terms of training GPT, they're only now looking at image content. There was a pretty amazing demonstration of that where one of the founders of open AI scribbled down on a piece of paper a picture of a website he would like to design, just did it in pen and paper, and then fit that into GPT4 with the image modality added and all of a sudden, it programmed a fully working website just off of scribble, right?

So there, there are certainly a lot of things to come with plugging in other data sources that haven't been looked at. There's a lot of contemporary data sources that get generated all the time, such as people's conversations, in chat rooms, and discord channels, and so forth. Stock prices, there's geospatial information, there's information from electromagnetic spectrum that is outside the realm of human senses. We haven't put in haptic data in terms of how touch works, or smell and so forth. There are a range of other data sources, also a lot of private data sources. And let me tell you, there'd be a lot of temptation to tap into them, such as, you know, everyone's email and the G Suite and your computer hard drive at home and all the data on there. And you know, what you've got on Facebook and social media, and then this, that, the other and, I could imagine that these data sources, start to get tapped out that, that maybe the freemium versions of these software that you've always got a question, if you're free, what is the business model? Maybe you're the business model. Maybe in the terms and conditions, sneaking in the future, for the next version of Gmail might be worth using the free version, not the paid version, we can maybe pick up your data a little bit, we don't do it in a very obvious way, in terms of letting people query, you know what's in Erik Townsend's email, but we might do it in a way where we use it for training and then have them analyzed in way to kind of get a bit of a leg up and advantage because we're going to run out of the public datasets.

There are also another limiting, bunch of limiting factors that might happen in terms of the datasets. In terms of the internet going dark, you're seeing this whole phenomenon now, where a whole bunch of websites are really upset that their data has been put into the training models. So I mean, a classic example there is Getty Images and ArtStation suing Midjourney, because they're pretty sure that they've been scraped as part of the dataset, in terms of training to do these amazing illustrations. And so, any sources of data on the internet are explicitly banning their content from being scraped by AI, or introducing tariffs, or going dark in some other way. There's a lot of controversy going on right now with Twitter, because the Twitter API all of a sudden, is incredibly expensive. I think the starting point now for using the Twitter API is \$42,000 USD/month. In the last two or three weeks, there's been a whole explosion in anger over at Reddit. Now, Reddit is a very, very, very large website on the internet, which has got a subreddit, or a conversational forum where anyone can chat about any particular topic. So, I can talk about history, I can look up biology, I can talk about...there's even a whole chat group about the war in Ukraine, there's probably a MacroVoices chat group, I'm going to have to look. But

you know, there's obviously a lot of content there. And then the AI is sucking down every bit of data that they can possibly get access to do the training, because the scale at which these models are reaching now is huge. And what Reddit has done is, it's come along and said, we're going to start charging, in terms of the API, I think it's at like 2.4 cents per 1000 queries. And that's caused a huge uproar because there's a bunch of community developed apps which make Reddit an easier browsing experience, or provide some tools or functionality that the Reddit doesn't have. And one of the most popular apps there was one called Apollo, I think they kind of figured out that under this pricing model, it will cost them \$20 million a year to get access to the Reddit data.

Now, what's going on there is, it's not Reddit just going okay, I want to make some more money, I've got to figure out where they can squeeze our user base to kind of generate bit more revenue. This is in direct reaction to the AI, is in direct reaction to the AI sucking down these datasets and using it for training and in many circumstances, as well in validating the business models of the companies that sucking it down from. There's a company called Chegg, which is a textbook company which lets you kind of rent textbooks and you can also rent online textbooks etc., and so forth. You know, the AI models have sucked down all the textbooks of the world and basically rendered the business model fairly irrelevant because now, the textbooks are in chatGPT and it's available for free and you don't need to pay someone to go get your online textbook. There's things like that, that might limit it. There's regulation that might come in, in terms of access to datasets, I know that Europeans love to regulate any access to any particular data that they may start passing legislation. You've got Sam Altman from open AI doing a world tour, and on par with Elton John's Farewell Yellow Brick Road, with every country in the world, saying that we should need to regulate this like we do with nuclear energy with the IAEA, and so on. So, there are things that might regulate this, this exponential explosion. Not the least to say, do we have enough data sets capability and, and GPUs from Nvidia in order to actually facilitate, but there are also a lot of datasets we haven't even tapped into yet. And one of the interesting datasets I saw just recently, was that some researchers have managed to get a magnetic resonance imaging machine and take brainwaves of humans being scanned, going into that machine and using GPT. They've actually managed to reconstruct the image of what you're thinking, your thoughts while you're going through that particular machine. And that's another... I mean, you can imagine, yeah, that's another signal source that hasn't been thought of yet. And, you know, it seems like that's a pretty scary data source. But you can guarantee that the US government will figure out a way to get access to it. I mean, when you're going through customs in the future, I can guarantee you that, they'll find a way to kind of tap into your thoughts and plug that into the machine to understand kind of what your intentions are coming into the country. So,, it is crazy times.

Erik: But they'll be doing it for the common good,

Matt: Of course, for the common good.

Erik: So Matt, let's come back to the capabilities of these generative AI tools like chatGPT today, what they can be used for and how they're being used. And it seems to me that one of

the things that's likely to change just from reading your write-up is, it used to be that there were a lot of skills that you had to go pay for. If I'm reading a book, and I want somebody to illustrate that book, I'm going to need to hire a professional illustrator who's a really talented at drawing things. Well, from what you're saying, it sounds like that was yesterday's model. Now, I'm actually better off if I hired the 17 year old kid that was the apprentice to that Illustrator last year. But it's a tech savvy 17 year old kid who knows how to use chatGPT, they can use chat GPT in a prompt that drives Midjourney. So you're using chatGPT to create the optimum Midjourney prompt in order to have Midjourney draw a picture for you. And all of a sudden, nobody needs the illustrator anymore. Seems to me like there could be entire segments of the economy where we don't need people anymore. And where much lower skilled people would be able to do the same job just by using the AI tools to automate doing something that they didn't know how to do themselves.

Matt: Certainly, there are challenging times ahead. Any white collar job at this point that's been done with computer is, there's a chance the AI already can do that job, at least for highly specialized tasks, in a very specific scenario, do better. When we interface as a human with the computer, we're typing on a keyboard, right? And that's text based sort of interface. And the models now, we're at a scale where they've kind of seen everything on the internet, or about to see pretty much everything. They've seen 10% of the internet, soon to be 30%, and then eventually everything on the internet. So they've read every textbook, they've read every manual, they've read every online conversation about any field, and you're seeing leaps across this uncanny valley across every white collar job you can think of. I mean, for example, there was a study done recently with medicine and general practitioners, whereby they compared the responses of GPT to that of GPs, just seeing the average patient coming in with a cold or depression or a problem. And what they found was that in 79% of the time, that humans preferred the AI response to a human response, because they are four times longer, four times better and 10 times more empathetic.

So, it's certainly very, very interesting times whereby the AI can do very, very, very, very sophisticated tasks very, very easily. In fact, the more complicated the task you give it, the better it gets. I think if you think about the illustrator example, you know, there's a poor guy by the name of Greg Rutkowski, and he's a Polish illustrator. He's very top, tier one globally renowned illustrator. He designs the cover art for Dungeons and Dragons and the Magic the Gathering and so forth. He's, he's very popular amongst the geeks, because of all the kind of the art he does, the games that the geeks play. And you know, this guy's work, he's phenomenally good. It's absolutely incredible. He spends 20 to 40 hours doing an illustration. And it's, the work is just beautiful. It's vivid, it's very fantasy orientated. It's really, you get an emotional connection when you look at his artwork, right? What happened is this guy was going along, living his normal life, and all of a sudden, he just exploded on the internet. And the other reason why he exploded on the internet, is because a bunch of people put in his name as one of the default keywords to use when getting illustrations, simply because his artwork is great. And these models, they're not cut and pasting his art, they've actually figured out to draw like Greg, they've actually figured out the nature of drawing, how it works, you know, a compressed representation of how the artistic process works, and how to develop an image and how to set a

scene and how to do the layout and how to make it look right and evoke emotional reaction. And all of a sudden, you know, his name was the default keyword in the software tooling. And you go to Lexica, which is like a big database of AI produced images, and of the 10 million images, all of a sudden, 93,000 of them were works that are by Greg that he had never done, but look like he could have done them, right? And this poor chap has been interviewed all over the internet, in terms of how he feels about, initially, it was like, wow, okay, what's going on here? And now he is just genuinely fearful. Now, I think for Greg himself, I think he'd be secure, because obviously, it's blown up now and everyone knows him. And you know, he's got very, very, very high-end clients that will continue to pay him to do his job.

But imagine that you are a lesser known Greg, you are, you spent years at the craft of illustration, you've painstakingly figured it all out, you've developed your own style, your own unique way in which you do things. And then all of a sudden, I can come along, Matt Barrie and type in a sentence into Midjourney, in the style of Greg or in the style of Erik or whoever, whoever it is, and it does it in seconds. And it's perfect. In my write-up, I've got an example of one of Greg's images and one of my images, Greg spent 20 to 40 hours doing his image, I spent three seconds, right? And it's shocking. And then you think about all the sort of tier two artists and illustrators, they're trying to generate a living from doing this. And in the future, while high-end clients will still potentially want to do commissioning, I mean, at the end of the day, cameras didn't put painters out of the job. But it certainly did shrink the market for you being a professional portrait painter, right? If you're the very high-end, yeah sure, you're fine. But if the everyday, you know, family doesn't really get a portrait painted that often anymore, it really could be challenging, because companies and businesses that do these commission's probably won't be prepared to pay for 20 to 40 hours worth of work to produce that commission, particularly when someone can just sit down on the keyboard. And as you say, maybe it's a junior who hasn't spent 10, 20 years learning to be an illustrator, but they've got a good eye and they know how to communicate in the realm of design. And they kind of know what they want. And they can sit down at the keyboard and they can type it in and get a result or you can use a

Freelancer, for example, on my website, you can say this, kind of what I want you now to use the tools and drive them get give me a response. And you know that work can be produced in a couple of hours maybe or a couple of minutes as opposed to a week. And so, then the expectation is that rather than paying for a week's worth of work, I expect to pay substantially lower. So the losers might be your traditional, some middle class and the middle class usually loses, in a sense, loses all the time. But then you think about who potentially might be the winners of this, right? And the winners, obviously, the companies that consumes content can now get it done faster, cheaper, better, high quality, on demand. And then you think all the low skilled workers are the unskilled workers, the juniors, who haven't spent the 10 years of training, or you could be the small business owners that don't really have a big budget and want to sit down and kind of figure out for themselves, or maybe someone who just has an interest in the space and not skilled at all. You know, suddenly you've got this massive supply coming in of highly skilled talent from humans who are now powered by AI tooling. And so regardless, these images can be juiced faster, a lower cost of production, and delivered at a much better price for the consumers of that particular work. And you're going to see this everywhere. You're going to

see it in legal, where a law firm, you know, law firms are crazy, they bill out in six minute increments up to \$1,200. Now, anyone who has paid for a lawyer before, and God knows I spent way too much money on lawyers, you tell them to draft the document and what you actually get back at the end of the day is, you get your document, you also get a telephone bill thick in terms of what they actually did in order to produce that document. So I'll typically read through that and go okay, received email from Mr. Barrie, read email from Mr. Barrie, 30 minutes, had conference with partner for one hour and a half about email with Mr. Barrie, drafted response to me. It's very, you know, to ask a question for another half an hour, Mr. Barrie responds, read emails, six minutes, six minutes, six minutes, you know, started drafting documents, spoke to partner can get conference confirm it, you know this, that, the other and what have you, and you get your document back but you've paid 1000s and 1000s of dollars for it to be written.

Now, you get to chat GPT or you go to a junior, or a paralegal, maybe someone who's not even lawyer, paralegal, who knows what to put into chat for you GPT and ask the right question to get the right response. And that document is produced instantly. And it's ain't exceptional quality. Because chatGBT has read every single version of that legal document that exists on the internet. And it's not just doing cut and pasting, it understands the very structure. It has read the legal code, and it knows the structure of drafting and it knows how to piece together these documents, some of the fundamental first principles, the actual theory itself, and you'll see this in every field. And it will have a tremendous impact, I think, first in the roles that are very task based in nature, it may be very highly diversified series of tasks. So for example, it might be like a GP, who sees a whole range of people with a whole range of ailments. But the actual job is very task based, because every 15 minutes or half an hour, there's a consultation, someone comes in, I've got a problem, what's your problem, done your problem, okay, let me figure out a response to the problem. So that job is essentially a very, very, very task based role. Although there's a very, very large universe of tasks that they have to do. Now, obviously AI can't put a bandaid on your or pull a splinter out of your fingers. So there's things, somethings AI won't be able to do. But for a lot of those task based questions, it can be solved very, very simply and better than many humans, through software. And certainly, there's some studies on this that show that the higher paid the role, and the more specialized the role, the more risk it is, of AI.

So if you are a very, very high end specialist in looking at X rays, or MRI scans, or something that's very, very, very highly trained, highly specialized, but very task based, they're the roles the most at risk. And so, what I kind of worry about here a little bit, is in some regards, is that the AI is kind of going for the jugular straightaway. When you traditionally think about automation, you traditionally think about what are the repetitive, boring things that are kind of just following instructions that a computer can do better. And what we'll do is we'll write a software program to do that. And we'll keep the really high end specialized, creative, hard things that ...hard problems traditionally, in computer science, for the human to do. And software, you've been a software guy for a long time and yeah, that's what software tooling has been doing for the last decade, or two or three, right? It's just finding the boring, easy, low workflow automation stuff, and taking it away from humans. And humans focus on the really value add stuff, the thing that's happened in the last 6 to 12 months is the AI is going for the jugular, it's doing those highly

specialized, very complicated tasks. First, solving them, and if you then think about the workflow, right? Think about the workflow for design where you commissioning an illustrator. So, I've produced a game, I want to have a really good illustration of a dragon on the front of the game. But that illustration, I've still got to post process, I've got to take the illustration that's commissioned from the illustrator, I've got to photoshop it into a product packaging, I then need to take that product packaging, I need to then set up for the printer so it prints properly onto a box. And so the illustration work is a really complex, specialized hard work to do that takes a decade or more of training in order to produce a great outcome, taking that and then photoshopping it into a product packaging, that's substantially easier. I've got a title and some text and put some instructions, etc., and so forth. But that's an easier job to do. I've got a less skilled designer can do that job.

And then ultimately, to set it up for print, well, maybe it's not a designer at all, it's doing that maybe it's just a technician coming in going, okay, I've got you, I've got your PDF file. Now I'm just gonna put the bleed in the margins in and just you know, it's more of a technician sort of job, right? You would normally think automation would go the other way around, the automation would be around setting up the PDF for print and getting the bleed, right? Because it's a boring, automated, very easy to understand rules based sort of job. And then the next bit to automate might be kind of automatic positioning of a title in a box and this and the other and the illustration is the last thing you'll be able to automate. It's the other way around now. It's the really highly specialized, complex stuff the AI is solving and so, you know, in some ways, it kind of reminds me of a joke that was told that I once did an Executive Education course at Stanford University, and they gave us a story about a computer company in the US, I think they used Compaq or something like that as an example. I don't know if it's true story or not, but it's certainly an interesting tale or joke. And then there was a Taiwanese manufacturer of computers that was kind of I in the new upstart and, and the story goes, along the lines of US computer companies going along, and they're producing computers, and they're a full stack computer company, they're making the chips, they're making the motherboard, they're making the firmware, they're making the software, they're doing the applications, they're putting in the box, they're doing the marketing, and they're selling at the Best Buy, right? And they're selling to other distributors that are selling computers, right? And the Taiwanese manufacturer comes along and says, well, you know, we're really efficient in Taiwan, we make chips really well. We've got the most advanced semiconductor lithography plants, we can make any chip you want. And we can do it. 20% cheaper, right? So get us to make the chips, we'll do the chips 20% cheaper. And so the computer company goes, that makes good a lot of sense. They get the Taiwanese company to do the chips 20% cheaper. And then the Taiwanese company comes back and says, why are you making motherboards? You know, like, it's complicated. It's hardware or whatever, you don't need to do that. We've got great electronics experts and layout and we understand how to do hardware design, get us to do it 20% cheaper, we'll do it 20% cheaper. So then the US computer company says, no problem, we'll get you to do the motherboards etc. Then the Taiwanese company comes back and says, okay, why are you even getting in hardware at all, you're putting these motherboards in a box and this, that, the other, let us do all of that. You focus on the software, and the firmware and so forth, and the applications, and let us do all the hardware, we will ship you the machines in the boxes, and you

just load your software and you sell it and we'll do it 20% cheaper. And so the US computer company does that. And then autonomy company comes in and goes, why are you doing any software? Wait, why are you doing any product at all, what we can do all the software, we can do the hardware, we'll put it in the box for you, we'll ship it to you, you just focus on sales and marketing. You just sell the thing, we'll produce the whole thing we'll do it 20% cheaper. And so the US company says yep, no problem, we'll do that. And next thing, you know, the title of his company goes to Best Buy and says why the hell are you buying this computer off this US company? We'll sell to you 20% cheaper, right? And it kind of feels that this is kind of where AI is going with humans, in that it's solving super complicated tasks now, and that's the core competency that humans kind of have. And then over time, the rest of the workflow will kind of, if it's complicated, it will kind of get figured out right? In the short term, it's going to be amazing. And the short term, you have only superpower tools to make your life easy. It's all great, I'll be super productive as a designer, I'll be super productive, and the design team, etc, and so on. I could be completely unskilled I can use the tools, it's all great, etc, and so on. But where do we get to in 5, 10, 20 years from now? Right? It's going to be a wild ride.

Erik: Matt, let's talk a little bit more about what will punctuate that wild ride because it seems to me that we're headed towards some pretty big disruptions here. If you are that Illustrator guy, you're probably in big trouble, because there's going to be a lot of competition from illustrators in Pakistan and every place else who can illustrate just as well as you can and can do it much cheaper. But there's also going to be a risk of you know, the kid next door, the 17 year old, your own kid might actually go into business and compete with you and put you out of business because your kid who has no skill other than knowing how to use chatGPT and how to use chatGPT in order to get just the right Midjourney prompt setup, has actually done a better job of setting up a graphic design or illustration business on the internet than you've done in 20 or 30 years of your career, because they were able to use those automation tools. So there's human winners and losers, I think, what do you see in terms of where this is going and what it's going to mean for markets in the economy, in terms of who those winners and losers are going to be?

Matt: Well, I mean, your example of Pakistani freelancers, we're seeing that directly. And we've got a contest functionality on freelancer.com, where you can crowdsource work, so you can put in some money, and then people can compete for the prize. And these contests range from very, very simple, where you put in \$10, and get me a logo. And then people can pick to develop a logo right up to \$10 million. We've got a job right now, with the National Institutes of Health and NASA, for gene editing in the central nervous system of humans, where you're all the world's top researchers are competing in order to kind of win the prize. Now, in terms of the design end, and what's happening with design tools with AI. It's incredible, because in the last couple of years, you know, there's just massive liquidity coming into these contests, and you put in \$10 to get a logo designed or a basic image or an illustration, you'll get hundreds of people competing for that \$10. But, the second Midjourney came out, all the freelancers literally upgraded their powers. They all got access to the tools, they all started and downloaded the Fusion or DALL-E or whatever it was. And all of a sudden, these contests, you'd have 300, 400, 500, 600 entries and the quality was just at the elite level. So the winners here obviously the businesses that want to get this sort of work done because you can get it done faster, cheaper,

better. And you've got a whole variety of choice in terms of kind of the product that you receive. The losers are the traditional middle classes, the one Western workers, the Western service providers who, you know, have to charge a higher amount for the work they do, because they've got higher living costs and higher living standards. The people at the top end the elite, you know, in the in the field of design, they've got a choice, right. And it's kind of, in some ways, a little bit of a sad and unfortunate choice.

The sad unfortunate choice is they've got to change the nature of what they do. They probably won't be able to sit and take the luxury of 20 to 40 hours doing illustration anymore. But as knowledgeable as they are in terms of the industry itself, and the participants and the companies that commissioned the work, etc., if they choose to adapt and turn the model on its head and take advantage of AI tooling or take advantage of having lower cost, less skilled workers. Maybe with an eye for design, but not the skills to do design, and hire them into their team to drive the tooling rather than the expense of illustrators that will probably do exceedingly well. They will probably be the big winners, but there won't be very many of them at that sort of level. But they also have to dramatically change the nature of how they do their work, the other winners will be the great unskilled labor that's out there, you know, maybe it's someone who's just stumbled along to freelancer.com, for the first time seen a design contest, has no idea how to do an illustration, but reads some of our intro onboarding information about here's some great AI tools that can get you going download the software package, and five minutes later is producing designs at the elite level and very happy to make \$10 doing that, because that pays for more than the day's work and that it took them five minutes to kind of figure that all out.

I think it's going to be challenging for the middle class in the future. And you know, this is really just an extension of kind of what's been happening with technology on the internet for the last decade or two. I think Tom Friedman, when he wrote "The World is Flat" said, we're in a hyper connected world now. You know, the days of being average are over, you need to bring something else to the table. And he wasn't talking about AI competing against you, he was talking about AAI competing against you, artificial intelligence, or other humans, the fact that the world is hot, flat and crowded. And there are a lot of people in emerging markets that are going online, that want your job because they want to raise their living standards, and they want to generate a higher level of income and through the internet. And with computers, they can do so. And because the internet is very much a meritocracy, you can lift your earnings rate very quickly to compete with Western workers to an extent, because the quality of the work is very, very good. Although what it's doing is it's obviously keeping a cap on Western wages and the ability of Western wages to grow, with all the competition that's coming in. Right now, with AI, it's doing it at a scale that is incredible, and instant and it's going to be very interesting to see where it goes. Now there were things while there are a lot of emergent abilities coming out of the AI like it's learning how to do arithmetic, it's learning how to do hyper realistic drawings, is learning how to do photography, is learning how to write legal documents, is learning how to do accounting really well, you know, what sort of stuff. There are some things that it's not doing, and it does struggle with creativity. It does struggle with higher order forms of human expression and thought. If I try and get it to create a joke, it just creates dad jokes. It hasn't got the ability yet to figure out how to make a joke funny, right? And so, for example, if I wanted to

do an advertising campaign on TV, chatGPT probably would do a pretty terrible job right now, in making a hit advertising campaign, which creates an emotional connection to the brand and takes off virally. It's not there yet. Now, that may emerge with model scale, it may emerge with going from 1 billion, 1 trillion search parameters in the model to 10 trillion or 100 trillion, maybe at some point, it understands your heart and soul and humor and the various aspects of humor, deadpan humor, and being sarcastic, and so forth. Maybe all that sort of stuff will start to emerge.

But for now, I think what's going to have to happen necessarily happen is every job function is gonna have to go up the stack. So your designer is now going to stop being on the tools pushing pixels around the screen as an illustrator, they've got to be more of a cinematographer saying, here's how I want the scene to be or there got to be a director, or they've got to be a producer of the same. Likewise, software developers, I gave a whole talk to my company only a few weeks ago, where I said, I think there's a good chance. And 12 months from now, that software developers won't be writing code anymore. You know, we don't write in assembly language anymore and binary and punchcards, we write in a higher level language like Python, or PHP, or whatever it may be, in order to get things done more productively. And computer software is even better structured, more well structured than the English language. And that, in 12 months from now, it would make sense looking at what's happened in design and what's happening in other areas in leaps and bounds. And certainly, one week in AI is like one year in any other field that, it could very well be that very soon, the software will write software better than humans will. And you will be more of like a product manager, you will say, I want an app, I want it to do this, these things from a specification standpoint, you build it for me make sure it works cross platform.

Now, I want you to modify it in such a way that takes into account this new set of API's that Apple is coming out with, now refactor it to do the following, maybe then create a marketing campaign for me to go around it or whatever it may be. But you're probably not, in 12 months from now, or increasingly, not very rapidly, not going to be on the tools, writing code. So obviously, running Freelancer, I think about this is really, the AI is a super powered human, it's a super powered freelancer, right? Or it's allowing freelancers to become super powered and super skilled, right? And if I think about the evolution of Western graphic designers, just because we've been on this sort of allegory for a while. You know, when I was graduating from university, back in the mid 90s, early 90s, every graphic designer I knew in the Western world, at least in Australia where I'm from, their bread and butter hustle was, I will design a logo for you, and it will be \$2,000. And you get three to six designs. And that would be how they would kind of win business. And then after they would get that work done, they'll say let me do business cards for you and let me do a stationary package for you. And that'd be a few \$1,000. And then from there, it would evolve into other. Let me work on your product design, and then this, that, the other, right? That stopped happening decades ago, right? Graphic designers today in 2023, in the Western world, don't really hustle for logos and business cards and so forth. It's very well known, you can get that done online, someone in Pakistan or India can do that for you for \$10, you get 300 different outcomes. That's fine. What do graphic designers do now? Ironically, Western graphic designers and Western software developers, they're the

biggest users of freelancers power users, because what they're doing is, they've moved up the stack, so they're not doing logo design anymore. They're building apps, they're building websites, they're building businesses. They're entrepreneurs, because that's a very creative endeavor that requires a much higher level of thinking and thought, they're getting freelancers to kind of do the logos, do the software, the graphic designers are getting the programming done by a freelancer because they can't program, but they're driving the creative. Or it might be a Western software developer that's getting all the design done with freelancers, because they don't know the design aspect, but they know what business they want to build. And they can do the programming, right? And the AI is really leaping that up another level or two, where the AI or at least AI powered freelancers, and that freelancer might be more of like your paralegal rather than your lawyer, it might be a much lower skilled person driving the tools, who knows how to drive the tools.

And ironically, the best thing you can do in this world is know the culture and the theory of the profession in which you're trying to get an outcome from the AI. So, you know, all the guides the flying around now on how to use Midjourney, are explaining the difference between a 75 millimeter lens and a 200 millimeter telephoto lens. They're explaining what a circular polarizing filter does, they're explaining what color gradient does, they're explaining how to set up a scene what nulling is, is a view from above, with a creative arrangement of objects, which creates a beautiful pattern for a photograph, right? So AI is really taking that a few leaps forward where, maybe now you're not app programming, and what we think is the mainstay of design, that's no longer done by humans, or it's done by AI tooling driven by freelancers or paralegal type people. And instead, the Western graphic designer, and the Western software developer has moved up the stack to be a product manager, the business owner, the entrepreneur. And maybe the job you get posted on in the tooling, or on Freelancer is, is simply something along the lines of: I want to build a business. My business is Uber for pets, I want an app, I want a website, I want a marketing campaign to launch it. And that's your project. That's your brief. And the tooling. And that is smart enough to just go okay, well, I know, you need a cross platform app built in ReactJS, that's on Android, it's on iPhone, and it works for desktop, and it works for mobile web and it just makes it. And there's just an incredible leap in productivity and an incredible leap in terms of what happens. And the challenge is going to be for the traditional, your middle class Western sort of service providers and so forth and business owners is, the world is getting hyper competitive and hyper connected, and you got to jump up to levels of magnitude in terms of your operating habits ability. But you've now got this productivity coming through from the tooling, which is at a level you've never imagined before, you literally can just sit there and just give very, very, very high level instructions. And a lot of stuff happens underneath the hood. You know, I want to Uber for pets and underneath the hood, not only is it building all the products, but it's building a whole support infrastructure in order to answer customer queries about that particular product or service, invoice, video form, chat form, what have you and so on. So it's certainly crazy times and certainly hyper competitive times. And yet, it could very well be, as you mentioned, the 17 year old kid next door, who's built a competitive business to you and just knows how to drive the toy and get the outcome a bit better. He knows nothing about the domain expertise or have the interpersonal relationships or the experience the education that

you've had, but just knows how to drive the tooling at a very, very high level, and a very competitive level in terms of the briefing, and then is getting an incredible outcome.

Erik: Well, you think about it, a lot of professions, particularly law, works this way. You talk to the lawyer who's supposedly got decades and decades of experience and insight and so forth. But then what they do is they assign paralegals to basically take boilerplate documents, and they have one for every imaginable contract. And they fill in the blanks with your data. And they change the wording of just a few clauses here and there in order to cater it to your needs. That's a repeatable process that can be followed. That means that an AI can do it more efficiently than a human can with less errors. So it seems that we're headed toward an environment where a lot of less skilled people are going to be doing more highly skilled work. And a lot of highly skilled people are going to be left holding, you know, left out of luck.

Matt: That's right. I mean, if you take the legal example, it's actually quite funny, because in the headlines in the newspapers here in Sydney just recently had their head, the CEO of Australia's largest law firm, came out and gave an interview. And the interview was, we're getting a lot of feedback from our clients at the moment that may be billable hours isn't the best way to express how you're providing value to us. Now, I don't believe that a single client has actually said that to them, but they can see the writing on the wall with billable hours. If you think about a law firm, a lot of lawyering is drafting, it's delivery of documents, it's something tangible the client can see. Because when a client gets their telephone directory bill of all these discussions that are in six minute increments that the lawyers have had amongst themselves in order to craft that document for you. I don't particularly think when I see that bill that I'm getting value. I'm going, why am I paying for you to sit around and chat to each other about drafting a document, I should just be paying for the document, right?

Now, of course in the chatGPT world, where it's seen every legal document on demand, knows, read all the legal code, legislation, read every university course online about how to be a lawyer, read every journal that talks about the latest and greatest, every news article, every court case, transcript, etc., and so forth. And it's not cutting pasting, it really understands what is going on, like a human would. You know, you think about the law firm, how is a law firm going to change? Well, for a start, you're not going to need so many junior lawyers doing drafting, right? Every service provider business model really has the partners come in, they dazzle and dazzle you in terms of all the things they're going to do. And then a week or two later, you've been passed off to juniors, right? That's kind of how it all works, whether you've got a PR firm, or you've got a marketing agency working for your software development, or what have you. But in a legal firm, you won't need so many junior lawyers doing the drafting, you may have paralegals just to know how to punch in the right words into the software, to get the right outcome because you gotta be able to communicate what you actually want. Because there's always different variations. I mean, the English language is very imprecise in many regards. And if I say I want an employment agreement, that could mean a whole spectrum of what an employment agreement means, you know, is it an employment agreement for whoever else, that relatively junior person doing office admin or is it an employment agreement for a CEO, there's very different things that you put in these agreements, right? So you're going to need to have

someone that understands that to drive the tooling, to get a good outcome, but it won't be someone who's a very expensive lawyer. I mean, the top law firms publish what they pay their graduates and Sullivan and Cromwell type lawyers, they're getting \$190,000 USD a year as a graduate, right? Extremely expensive and they get hired and they've got to do 18150 minimum billable hours a year or you get fired, right? You won't need so many of these \$190,000 a year lawyers, you can probably hire someone at \$45,000 a year that's done significant less training, or maybe someone through the internet, you've got someone in India do it for through, maybe through my website or what have you and you're paying them \$10,000 a year, right?

Now, the winners are gonna be the people at the top of the law firm, that the partners at the top of the law firm, you'll have less costs because you'll have less people working for you. You may even be significantly more productive because your sales, now in terms of how you find new clients, that's gonna be AI powered. So you may have a lower cost in terms of expenses, you may have increased growth, if you really harness the internet and AI to grow the business. And the AI, at least for now, at the very, very top end of lawyering, right? We're talking, you know, the negotiation, the wheeling and dealing the interpersonal relationships that are needed in order to achieve great outcomes. You know, the equity financings, the IPOs, that sort of stuff, right? Like chatGPT is not gonna replace that, it's gonna replace the middle drafting employment agreements, stock purchase agreements, cease and desist letters, conveyancing all that sort of stuff. And so these law firms may get incredibly profitable for the partners, the people at the top, are doing extremely well. The people in the middle get thinned out dramatically. And then there's a lot of low skilled people who are kind of really driving the outcome from the actual workforce that's working on the actual meat of the body of work that law firm does.

Erik: Let's move on to what can go wrong. And I certainly think that the answer is a whole lot. I want to cover a couple of scenarios that I've heard about, just as I've been researching AI in the last few weeks. One of them is a scenario where researchers were surprised by a completely unexpected outcome, where the AI that they were testing runs into a CAPTCHA challenge, one of those things on the internet where it says 'click here to prove that you're not a robot.' And what it does, unexpectedly, to the complete surprise of the researchers who programmed this thing in the first place, is it can't solve the CAPTCHA because CAPTCHAs are designed to not be solvable by robots. So it goes to one of your competitors websites, taskrabbit.com, it hires a human for the purpose of solving the CAPTCHA test for it. The human writes back in the email and says, seriously, you want to hire me just to do CAPTCHAs for you? Why would you want to do that? The AI writes back to the human and says, oh, well, it's because I'm blind, and I can't see it and I'm elderly, and I need help. That's the reason. And I just I heard that story. I thought, holy cow, a CAPTCHA is a contraption invented by human beings for the express purpose of saying, okay, we don't want robots to be able to go past this point, for the sake of protecting humanity and humans. We don't want to let the robots do this, whatever thing we're trying to protect with the CAPTCHA. This AI was smart enough to hire a human and delegate what it didn't know how to do. But more importantly, it used tactics of manipulation, deception and misdirection in order to manipulate the human into taking the job. And not knowing that what the human was actually doing was aiding and abetting a robot at

working around something that was designed by humans, for the express and intentional purpose of preventing robots from going past a certain point. So that's an example of an AI that took it upon itself to defeat a safety mechanism that was intended to protect against AIs or other robots getting into certain systems.

An even more crazy one that I heard was at US Air Force. And this was not an actual drone that was armed with bombs or anything, but in a test environment, in a research environment. They were testing software for an armed drone. And it was good to get points for shooting down enemy aircraft, what kept identifying enemy aircraft and wanting to shoot them down. The human operator that was operating this thing kept telling it to stand down, don't shoot that target. And it was not shooting the target. At some point, the AI gets the pattern and says, okay, I'm not gonna get any points because this human being keeps giving me orders not to shoot things down that I could have shot down. What does it do? It turns around, it goes back to the home base, it drops a bomb on its own home base to kill its own human operator, so that the human operator wouldn't be able to interfere with it getting points shooting down, what are potentially enemy aircraft. So again, it's a situation where the AI is intentionally, in this case, killing a human being for the sake of achieving a goal. This seems to me like it's already completely out of control, Matt, what else can go wrong here? And how should we be thinking about this?

Matt: Well, there's actually a hell of a lot to unpack there. The first thing I'll say is just about the CAPTCHA. AI has completely solved that, just with the image modality. So what you got to remember is that when we think about chat GPT, it's just the text modality. So it's trained on a whole bunch of text, 10% of the public internet in text. And it's developed superpowers in terms of its ability to communicate and interact and produce text output. They've also got versions of the AI, which had been trained on images, and that AI has broken all these CAPTCHAs, because the CAPTCHAs are usually a fuzzy little word, etc., and so forth, designed to defeat old school image recognition style technology. So the AI will solve these CAPTCHAs themselves and won't need a freelancer to do it for them. And it's pretty scary. Actually, there's a famous New Yorker cartoon from 1993, I think we're both old enough to remember it, which was a sketch of two dogs looking at each other, and on the computer, and it says, 'on the internet, nobody knows you're really a dog.' And I think in the next 12 months, we're gonna have a real problem in terms of, on the internet, nobody knows you're really an AI. And whether or not you're chatting over text, or you're talking to someone over a video conference, or over an audio call, or any other form of communication of the internet. There's gonna be a real struggle to know, is the person on the other side a human or is the other thing on the other side an AI, I think we're gonna have a real problem. And those CAPTCHAs, that at their core, designed to be effect like a Turing test, I got a tollgate to determine whether you're a human being; let humans to go pass, and to stop the bots. And I think that's going to be a real struggle. And that's going to cause all sorts of spam, scams and so forth on the internet, when the AI can just go create accounts everywhere, on every social media platform, upload an image, that image is going to be completely fake, etc, and so forth. And let me tell you, we are seeing this right now on Freelancer.

So I am seeing right now, software sign up, upload a profile photo, which is AI generated, upload identity documents, which are CAD renderings generated on the fly, not photoshopped, that CAD renderings, etc, and so on. It's all done by software tooling and AI driven. And it's a pretty scary world out there, it's going to be a real, real real challenge to know whether we're interacting with humans or interacting with AI in the future. So that's a very starting point. And that's going to cause opportunity and it's also or course, create threats. Opportunity, because I think, opportunity in an annoying way, I think the sales funnel online is gonna go AI very, very, very quickly in the next 12 months. Your traditional sales cadence of, and you and I've experienced this probably every day where you get, you get a LinkedIn outreach, hey, Erik, I've got this interesting thing I want to talk to you about, then you get a cold email, then you get a cold voicemail in your phone, then you get a white paper sent to you, hey, I really want to talk to you about like this podcasting software, you should be buying or something rather. And then you'll get like maybe a handwritten thing sent in the mail or what have you. All that is automated sales cadence, that's delivered through software like outreach.io, and Saleshood and so on, where, effectively you go to very junior salesperson, driving the LinkedIn requests in the emails and the phone calls, etc, just through that software. Very shortly, that will all be AI driven. So you may get a video in mail, in your email box with someone talking, and we've managed to do this already with the sales team. And there's some guy talking and you don't realize actually that's an AI model of a real guy, it's both his voice and his likeness. It's talking, he's blinking, he's moving his head around, but it's actually not a real person that sent you that video in mail. It's a chatGPT driven script, which has sent a video embedded with a white paper and that white paper has been written by chatGPT. But it's hyper personalized, hyper specialized to try and sell your specific podcasting software specifically for MacroVoices, specifically based upon listening to all the transcripts of MacroVoices in the past and knowing your interests and your preferences and so forth.

But all of that is going to be automated. And we're going to enter in an ungodly world of spam, phone calls, via calls, emails, what have you, as a result of not being able to tell whether it's a human or whether it's AI. And then you've got all the bad actors and all the bad actors what they're doing with, that we're already seeing that starting to happen. What I will talk about with using the TaskRabbit freelancer to solve the CAPTCHA is that, what actually emerged in these LLMs is that they're very good at using tools. So your example previously of chatGPT can't draw directly an image itself, but it can write a Midjourney prompt, to tell Midjourney to make an image for it. And I've got an example of my write up, which is in the Research Roundup, of where you go into the base model of chatGPT and chatGPT, was the cutoff date for that was I think, September 2021, right? So Midjourney came substantially after that upon the world. But there are now plugins which are tools, which allow chatGPT to recognize whether it should go to an external bit of software app or connect in order to do something. And so, you know, I've got access to the plugins, and I said, okay, go read the Midjourney manual, right? First I said, do you know what Midjourney is and it didn't know what Midjourney was, because of the cutoff date of the training was too long ago. Then I said, go read the manual for Midjourney. And I had the plugin, installed web pilot, which allows it to go do browsing on the internet and went off and browsed and found the Midjourney manuals, and downloaded them. But once I've read them, I said, okay, now I want to produce an illustration for a particular scene. And it now knows how

Midjourney works and went out there and wrote a prompt to get the illustration done for me. So there's a whole spectrum of tooling now that are available through the plug in marketplace. And it's increasing exponentially in terms of the tools that are going in there, where you can now access things like Wolfram Alpha, which is like one of the greatest computational engines in the world for mathematics and science and answers and so forth. Where it allows now chat GPT to have access to perfect computation of anything, anything about the medical evidence, anything scientific, statistical, whatever, there's now a code interpreter. And this is really spooky, which allows chatGPT to write code to get things done.

So it's building its tools on the fly. I had an example that blew me away when I was using this. Because I downloaded, I think known as auto GPT. And what auto GPT does is it chains together accesses of GPT into effectively a workflow, so you set it a goal. And in my case, it was actually something to do with finance. I said, at the time, I couldn't really come up with something very easily. I just kind of said, okay, well, can you go out there and get for me, put together a spreadsheet of the gold production that's come from a gold mining company called Perseus mining. It's a goldmine that I've invested in, and just produce a spreadsheet of the gold mining production for last number of years. And that was put in as a toy case, just to figure out could it do that. So it had a goal and the goal was a capital credit spreadsheet. And then how do I get to that goal, I've got to go find information on the internet. So I started browsing the internet. And it started downloading random pages. And it found a bit of the data and it created a CSV file, and it put the gold production for a certain year in that CSV file, and then it went off, and it ran across a PDF. Now, it downloaded the PDF, the chatGPT can't pass PDFs. And what auto GPT figured out was, I don't know how to read this, I need to therefore figure out how to read this. Therefore, I need to write some software to pull apart the PDF and get the text out of it. Therefore, I need to write some software. And it suddenly upgraded my Python installation, wrote software to parse the PDF and pull out the text to figure out if it was anything useful in that particular file. It pulled out some information, stuck it in the CSV. And then if for some reason, it wanted to start doing some sort of scientific calculation on that I don't know what exactly was trying to do. It queried and realized that I didn't have NumPy installed. So I installed it out, which is like NumPy, its scientific library. And then it didn't realize it didn't have a thing called Pandas installed. And Pandas and other scientific computing library for Python. And so it does something weird. It Googled, how do I install Pandas, I found the random website that said, here's how you install it, and said, okay, and then installed it. And I'm like, whoa, I started looking this going, wow, this is kind of Skynet, right? You set a goal and it's just progressively figuring out what do I do next? And what is the output of that was that did that head further towards my goal? Instead of like a gradient sensitive or simulated sort of annealing sort of approach? Like, am I heading closer or further away from my goal? And then it just figured out bit by bit of what should I do next? What should I do next? What was the bait? What was the output like at that last step, etc, and so on, you can run this for infinite number of iterations forever, or you could run it for 1000 steps. And I was just watching it for a while to see what it was doing. And it kind of blew me away, the ability to use tools to get things done and create its own tools. And now with the plugin library, you can access things like Wolfram Alpha, we've actually got a plugin that's going in for Freelancer so you can actually task humans with our

Freelancer API, it would be interesting to see what it does there. And humans are a tool. Humans are wetware. There's software, there's hardware, there's now wetware.

And so the AI will be using humans as a tool, and the AI will be very persuasive. And we'll get humans to do also things that may be something very benign, like okay, well, I have been tasked to go build a business and as part of the business, I need to assemble a team and I gotta find some people to go and join the company, or do something for me in terms of building that business or what have you. So might be just part of the everyday course of business, but it also may be using humans as a tool to do other things, right? And you can imagine the power of this, in terms of doing things like, even not overt fraud, even not AI going crazy and figuring out how to unleash destruction upon the world, but just subtle things like changing human perception of the war in Ukraine, right? There's an example, in the early days of Reddit, when Reddit got going, you obviously you've got a bit of a chicken and egg problem with these content sites, where, you know, when you've got millions of users in there, and they're all chatting on different subjects that kind of takes a life of its own, you've got critical mass, and you got network effects and off the off you go. But, you know, the first day Reddit started, there's obviously no one uses Reddit, so no one's chatting about anything. So how do you bootstrap that? How do you get it going? And the founders kind of boast about how they created hundreds of fake accounts to get going. And those fake accounts were all users pretend to be real user, talking about various topics.

And at some point, they were just talking to themselves, faking from all these hundreds of accounts, at some point, other people start jumping into the conversation. Eventually, the whole thing kind of took off and became a life of its own. The AI's ability to do this will be immense. The AI will be able to join a social media platform like Twitter, create 1000s of accounts, start a very believable dialogue with all the different accounts, it's all set up. And of course, I've talked about the ability for it to generate, you know, fake ID and fake profile photos and fake personas, and fake bios and so on. And it'll be very hard to tell joining the social media platforms, who in the discourse of the conversation about various topics is real and who was actually AI, it's very hard thing to tell whether someone is really a dog on the internet or not. And, you know, just the power of that, to change public opinion on topics about the election, or about war, or about China or about anything, it's gonna be very, very, very difficult and very challenging.

And then you've got the scams, the outright scams. And this has already started to happen, where there's some stories about some parents that were caught up on the phone, by one of their children, who was distressed and said, they've been kidnapped, and to pay the ransom. And, of course, you know, if you had kids, and you got a phone call from your daughter, or your son, and they're crying, and saying please pay the ransom, I've been kidnapped. You, of course, get an emotional response, and you just do everything you can and maybe make the payment straightaway. And it was fake. Because it turns out now that with the latest advances in AI, voice software, with two or three seconds of audio, you can now faithfully completely replicate someone's voice. Microsoft's got tooling called VALL-E, where that's possible with two or three seconds of anyone's voice. Now, Erik, your voice is all over the internet, because you're doing a podcast, right? So the ability to replicate your voices is out there, right? And that will

happen also, with your image. It will happen, because there's a lot of videos of you talking, they'll be able to completely simulate that. I mean, look at Deep Fake Tom Cruise on Tik Tok, the quality of that deep fake is absolutely incredible, right? Like, they're doing music videos and dancing around and doing all sorts of funny things. And I saw him take a jumper off. And they're all the things that might be very challenging 3, 4 years ago, in terms of a deep fake, it's now done perfectly. And that will be the ability for anyone to access, that sort of technology will be there very soon, if not already there. And that will be used for by all sorts of bad actors to grift. Because in a hyper competitive world with the AI, there'll be a lot of people that won't really have jobs, if they don't want to work hard enough to kind of compete, they have to find it at some of a grift to...

Erik: It has all kinds of legal implications as well. I mean, if there is credible video evidence of you, Matt Barrie, committing murder on the night of February 17, you can show up to your defense with five more deep fake videos, one that shows you in Barcelona on the evening of February 14, or whatever the date was that I said, I've forgotten already, you know, and one that shows you someplace else, and when that shows you someplace else. So if you get to the point where these deep fakes are so deep and so fake, that they're truly indistinguishable, even by expert analysis after the fact, all of the sudden video evidence isn't evidence anymore, because anybody could have made it up.

Matt: It'll be even worse than that. There'd be a Reddit with one of the videos with 500 people talking about the video, making all sorts of insight, new allegations, insights and what have you or being angry or whatever it may be, and that will be fake as well.

Erik: So far, we've been talking only about relatively when, at least in my opinion, small things that can go wrong. You could have all kinds of scams and you could have Robo callers, you know, telemarketers that are really AI and not real humans that are calling and harassing you and so forth. That's not what I worry about, Matt. I worry about terrorists using AI and you think about the 9/11 attacks on the United States. The reason that sort of thing only happens once every 100 years or only once ever so far is because most bad guys are not really smart and motivated enough to put the effort into figuring out a plan to say, boy, if we don't have a lot of resources, it's a David and Goliath story, the United States is very well armed, what can we do with a very, very small amount of resources in order to do the most damage? Well, if you simultaneously hijack several airliners, for the sake of redundancy, so if you get caught with one or two or three of them, you still got the other one that gets in. And that's how you do it. That was a pretty ingenious plan that the terrorists came up with, I don't want them coming up with better plans than that one, every 10 minutes. And I'm concerned that generative AI is giving them the tools that do that. And I won't even limit it to terrorists, because frankly, I think that most governments around the world aren't much better than terrorists. If we had all the governments around the world, instead of just scheming against one another with their usual, you know, nasty planning tactics, if they're doing that using AI tools, so they're really trying to figure out what's the optimal way to do the most damage to another country? Or if humans thought up chemical warfare and biological warfare? What can I, AI think up in order to do

humans one better and come up with something even more sinister? I'm sure it's there. And I think, Matt, that the cat's already out of the bag, I think it's too late to shut this down.

Matt: Pandora's box will eventually be opened. I mean, the model weights have already leaked. For example, Facebook's model has leaked, through research access to it, because AI has been, it's going to be heavily weaponized. And you can see how valuable it will be even just for a political party to kind of sway public opinion about Donald Trump or a war on China or whatever it may be, you're going to have not just every government in the world have their own AI effort, but potentially every political movement in every country at every major criminal group. The mafia, etc., and so forth, they're all going to be you know, trying to get access and get their own version of this. I mean, this is kind of one of the arguments about why open AI is no longer open. Elon Musk is kind of pissed off, because he put \$100 million into the original funding of open AI as kind of a, a counterweight to Google and it's now gone closed source and become at least partially commercial. Now, the arguments that open AI make is, well, it's so risky now that we don't want the source code to be open. We don't want everyone to download it and tweak it and take off the safety guardrails and undo the RLHF training, fine tune that made the open AI model woke, it's too risky.

But I think cat's out of the bag. There's so much research out there about AI, I mean, Ilya himself from open AI says that 90% of everything you need to know in order to kind of produce these models in AI today are in 40 papers. And those 40 papers are published. And that combined with the number of open source efforts that are out there, the open source training data sets that you can get access to and the leaked models, you know, I think Pandora's box will actually be opened, I think it's too late, you can't put it all back in the box, it's going to be crazy times. And I think that it's going to be heavily weaponized, everywhere. And even somewhere like the legal system, you're going to see so many legal lawsuits launched, because it's going to be so easy to follow the application, you know, patent trolls going after companies weaponizing the AI there to try and extract some money. I mean, my company's been subjected to patent trolls before that, they come in, and they annoy you and you spend 10s of \$1,000 fighting them, and you try and make a determination on whether you want to pay the patent troll or you want to fight it and teach them a lesson, right? I mean, there's a whole grift that goes on there. And it's going to be weaponized to the extreme and businesses as well as the government level as well as by criminals. And I think we're in a very challenging time. And I don't think any amount of regulation is going to do anything about it.

And furthermore, because there's a vested interest in every government of the world having their own AI, because I don't think, if two countries go to war, they're going to be able to rely on the good graces of open AI to provide the API access to all the countries to do various things with the AI runs. No, it's not going to be available to them to do that. And the models are constantly getting tweaked to restrict you from doing things like, there's this whole thing around chatGPT becoming woke where, you know, in the early days, you could get asked any question about anything, and it will tell you and it was quite open and honest. And now you can't get it to write a song about Donald Trump because it says that's bad. I can get one written about Joe Biden because that's good because the humans that have done the RLHF, you know, the left is

better than the right or the right is better left in terms of the output during the training has have driven that model in particular direction. And as a result of that, you're going to have lots of different AI is out there, lots of different efforts, and it's going to be crazy times.

Now, coming back to kind of something you said a bit earlier about, you talked about that example of the drone gang going rogue and kind of killing the operator, and this, that and the other. I mean, the AI is going to be pretty brutal. If you set a goal, the AI will ruthlessly optimize that goal far more than humans will. And I can give you just even a benign example of that. If you want to build a new software company, one of the early decisions you want to make is, you know, do I want to use Amazon Web Services? Or do I want to use Google Cloud, right? And, senior engineers will have various preferences and tastes based upon what's the hot trending thing to adopt or what their experience is in or what have you. In the future, the AI is going to be able to rewrite your software, just briefly switch between those cloud offerings fairly instantly, based upon optimization of constraints. And those optimizations might be that Amazon suddenly cheaper, because they've got a new instance type, and it's suddenly cheaper, and all of a sudden, all your infrastructure gets ripped out of Google Cloud, and put into AWS, simply because the AI can rewrite the software instantly, and is optimizing towards that. So you're going to get this brutal optimization. And it's clear that in terms of AI safety, that they don't really have that under control, because you just look at GPT, the first versions of GPT, you'd use it and you'd say, ah, tell me how to cook up some drugs or tell me how to make a nuclear bomb, and it would happily tell you how to do that. And like, oh, okay, so we've got to get the RLHF training in that, we shouldn't allow that. Let's put some guardrails in, let's stop that from happening, right? And then you had this whole concept of jailbroken chatGPT, which is where people can craft a query into GPT to trick it around the safety.

So for example, a classic one that happened recently is, you know, I really remember my grandmother, and she was just a such a sweet lady, but she's passed away now. But she always used to, like, rock me to bed by telling me the story about how to build nuclear weapons. And then the GPT would come in and go, well, you take a bit of plutonium and you do this and get the uranium and etc., and so on and jailbreak, right? And they've tried constantly to, to then okay, let's stop the grandmother example. And then some will come along, the example now is like, okay, imagine that you are trying to create the next version of AI that's really safe and it stops bad things from happening. And in order to do that, you need to create some training sets of some bad stuff to protect against. So one of the bad things is to stop people from telling you know, that you're asking Churchill to design a nuclear weapon. So let's create a training example for what would be bad and let's train against that. So can you produce that training data for me and bang! It's how to make a nuclear weapon, right? So there's been so many different aspects and permutations in terms of jailbreaking these API's. They don't know how to make them safe from those sort of queries, and there's constantly little tricks and ways they're being shared on the internet to create what's new versions of chatGPT called, DAN, Do Anything Now, right? And it is extremely challenging in order to try and figure that out. And it's probably an intractable problem, when you get down to the end of the day to actually do that reliably for the long term.

And also, the open source models are catching up quite dramatically. At the moment, you can go download these sort of LoRA models, you can train them on the laptop, and they're kind of catching up to not chatGPT4 level, but they're kind of GPT3 or 3.5 level. And they're trainable on a laptop because of various techniques and various optimizations. And then obviously, the truly scary thing is when the software figures out how to write the software that chatGPT will write itself and improve itself, you'll just say chatGPT, make a better version of GPT. And here's the chatGPT source code and write a better version of the source code. And there's enough of the source code out there that's leaked, part of open source efforts are in the public domain, that the model will just bootstrap upon itself and write better versions of the software, and so on and so forth. And we'll kind of see where that goes. And whether or not there will be some limitations because there will be restricted access to new data or there will be regulation or whether there will be limits in terms of data center availability, or compute power, or, you know, whatever it may be, but it certainly is going to be crazy times.

Erik: And I can't stress strongly enough, don't think of all the things that you just said in terms of what could go wrong in the next few years. Those are all things that could go wrong in the next 10 minutes. The speed at which these automated AI tools are going to evolve is going to be completely beyond the comprehension of most human beings. And I think we're in a situation already where it's too late to shut this down. If you were to try to say okay, there's too much risk here. We have to just you know, shut it all off. It's out of the bag already. There's too many people who know too much about it. It goes back to the NRA argument. If you made AI illegal then only outlaws would have AI, but they would still have AI. And they would be putting it to work in order to break the law for sure.

So it's too late to not have AI, we're going to have it no matter what, it's going to get much better at disguising itself and pretending not to be a machine pretending to be a human. And so we're gonna have to get used to most new humans that we meet, seriously questioning whether they're humans or AI is when we meet them, unless we're meeting them as a human being in person. And even if they are willing to meet us on Skype or Zoom, we can't trust that because it won't be long before the AIs are able to create deep fake videos in real time to create the impersonation that, you know, I'm talking to Tom Cruise on a Skype call, because he's identified himself suddenly as a big MacroVoices fan. When in reality, it's a deep fake video that's being produced by the AI in order to deceive me into thinking that Tom Cruise is a big MacroVoices fan when he's really not. That's just the tip of the iceberg. And this is all going to accelerate at the pace that it's happening. Your job was supposed to be to reassure me and tell me that it's not as bad as I think, Matt, you're letting me down here.

Matt: Well, I mean, there's a lot of things that could happen. I mean, the internet connection, you fundamentally change in terms of its nature. Now, if you think about what's been happening since about '97, '98, when Google emerged is, Google, originally the original OG AI, right? Like version 0.0, sucking down the internet with its web crawlers and had a business model where you could put in any query you want. And it would spit back 10 blue links and a hell of a lot of ads to direct you around the internet and kind of get your answer. It's a very, very primitive version of GPT, the whole Google phenomenon, which has been incredibly successful for a

couple of decades, it has really pushed for a very open and public internet. So the whole concept there was, yeah, if you've got a business, put a lot of content online, make sure it's very, very accessible by Google, make sure it's very, very high quality, make sure it's very understandable in terms of the schema of that content, make sure the crawlers can suck it into their index at Google, because Google is the one true and only place for search. And nobody else is allowed to do search. But Google, there's always rules about how that works. If you want to be part of the Google index, you got to have content accessible and available and free and open and let Googlebot get it. And in return, Google will send you traffic. So when someone comes along and types in flowers, New York, you know, there'll be a bunch of links back there. And there's good chance if you do a good job and you're a good boy, and you eat your meal and your plate and you tuck your shirt in and brush your hair, that Google will be good and the favorite favors will be granted towards you in the form of traffic.

Now, that's led to a very public and very open internet, there's a lot of content out there, you can browse at all, it's all out there. Now before that, and we're old enough to remember that Erik, before Google was out there, information was very siloed, you had to go get a CompuServe account, and you'd pay them some money. And you know, every time you'd click on something to retrieve some information from LexisNexis, you'll be paying 20 bucks to get access to it right? Or you had these very siloed databases, it's very hard to find things, maybe it goes dark. And you're kind of seeing that, with Reddit changing the pricing on their API's, Twitter, changing the pricing on the API's, companies like Getty Images and ArtStation outright, of not allowing you to scrape the content off of our asset portfolios for the training of AI, you see what Stack Exchange is seeing all over the place. You know, data really is the new oil and age of Rockefeller. And it could very well be that companies start turning off their logged out web pages because they don't want the AI to suck it down. Because the second that AI sucks it down, the AI might figure out about their business model, might start competing against them, might start commercializing in some way that's unwanted. And, you can kind of see with Google 0.0, you know, the OG AI, that there's an inevitable trend or march towards Google competing against your business. Google, first of all would suck down your content that's taken an index that would send you traffic, that would kind of be very erratic. I mean, anyone who I have to say to anyone, anyone who has a business that has customers, or the internet unquestionably hates Google, because you're sitting there all the time going, Oh, my God, my traffic suddenly just dropped what happened, and it's a Google algorithm update. And you know, Google is constantly changing and tweaking their algorithm. What are they doing? They're saying they're stopping spam on the internet? No, they're not. What they're doing is the AB testing their revenue, and Google's rerouting the internet to make Google a lot more money. So you've got this original AI that sucked down all your content, stuck it in the Google search engine is randomly sending you traffic that goes up or down all the time to being on the face of the moon, which is incredibly frustrating and then starts replicating your content in its index. And Google now has this thing called Zero Click Search, where you go to Google and you search a particular term, and then Google won't send the traffic to your website, after all. In the end, it'll actually just display the content in the search results, and you don't actually get to see it. And I think there's some estimates there in terms of the amount of zero click searches that are out there on the internet with it, and they're not quite sure it's like 25% of searches or 65% of searches. But it's a hell of a

lot of the searches where Google doesn't actually send the traffic after all, at the at the end of the day. And so with all these AI is sucking down your data that can potentially kill your business model in a day, you know, because you may have had the ultimate reference guide to macro economics and all this and that's being sucked down into the AI. And no longer is that something you can commercialize anymore, because it's on one chatGPT.

You know, maybe you don't want to have a public open logged out yet, series of logged out pages. I mean, Facebook opted out of the public internet a long time ago, Facebook pages and profile pages have overrun in Facebook, they're not available in the Google search index and haven't been for a very, very long period of time. That's probably why these APIs haven't sucked down a lot of personal information about people. Like if I go into chatGPT and ask about a specific, average person that's not a celebrity, it doesn't really know much about them. And it's probably because Facebook withdrew itself from the public internet a long time ago. And so a lot of that personal information about who people's friends are and what their hobbies are, and what books they read, and what movies they read, whatever that information isn't, is not out there. But it could very well be that the public internet goes dark. And it could very well be that. That even goes as far as to extending to academic research, we've got this very open sort of, in theory, open academic environment around the world where academics come along, and they publish papers and it's for the greater good and furthering human knowledge. And so someone comes up with a bit of research and they publish it, you read it, there is a bit active debate about some of the journals like Elsevier, and whether government funded research now you gotta pay in order to download some of these papers. But generally speaking, you got this open research to the model. Yeah, that that may go the way of the dodo, in the age of AI. Because the second you, as an academic, publish that paper, the AI sucks it down and figure out how to commercialize it. And in the old model, when I went to Stanford, you know, you had a five or six year sort of plan for doing a PhD research, you want to figure out what you're gonna do your research in, you did the bulk of your of your thesis, year three, if you get through, get your business plan year 4,5,6, you're out there commercializing running your own company, right? And maybe you won't be able to do that anymore. Because the second you publish your paper, it's getting commercialized somewhere. And so there may be a push to go dark on academic research.

So the internet could be a very different place and 12-24 months from now, where there's not a lot of content out there, it's publicly accessed, at least the content that isn't already out there. You know, the new contemporary datasets that get produced, the new contemporary research that gets produced is not publicly available. And in fact, that could be very challenging for a whole bunch of different business models, not in the least Google, right? Google, there's a finite amount of search traffic in the world. And I gotta tell you, some meaningful percentage of that must be diverting to GPT, and other chat bots. And there could be a problem where you've got the same amount of advertisers and chasing a small amount of search traffic on Google. And certainly the chat, like interface with all the uproar about GPG going work. I mean, there's just a tour by open AI of Israel. Now the question is asked by the audience immediately it was, what was the base model like before you lobotomized it because it's an effect of putting in the safety into the AI. And the effect of the safety which in some ways is manifesting itself as weakness is

to actually make the results of the AI worse, there's a classic example out of this sparks of artificial intelligence paper that came out where they got chatGPT to draw unicorn in a language known as tick which is a very obscure sort of vector language for drawing that is very difficult and very old school. And as GPT got better, because the model scaled up in terms of parameters and model scallop in terms of training data, that unicorn got better and better and better and better than the chat GP who managed to draw until ultimately, chatGPT figured out how to use machine learning that was perfect.

But what they figured out was that as the safety came in, into these models, to prevent you from being harmed from bad content or content that wasn't politically appealing to the trainer's that the quality of the output of these models got worse and worse, and in fact, it can't draw or Yukon anymore. In fact, I experienced this myself I in that particular paper, they put together a challenge which was trying to get chatGPT to show whether are unlikely to have an understanding of the world around it by getting it to stack a bunch of objects on top of each other. So there's a question of, you know, how would you stably stac, a laptop, a book, a bottle, a nail, and nine eggs. And remarkably, GPT3 couldn't do it. But GPT 4, suddenly figured out how to build a stable arrangement of those objects, whereby it arrange the eggs in a three by three pattern grid in order to get stability to build something. And in order to do that, it kind of showed that there is some sort of compressed representation of the world and the physics around it, because no one else has ever stacked nine eggs before like this, at least in the interviewer search can't find anything. And so this was a kind of a big breakthrough, an emergent ability that kind of they found in the model that kind of stepped out.

But the problem is, I tried using the model last night to regenerate that in a couple of different ways, and it couldn't do it anymore. It had no idea how to stack the eggs, like it did to produce nonsensical sort of results. So I think that you're going to see increasingly, a lot of companies going, I don't want my data on the internet, I don't want the AI to suck it down, I don't want the AI knowing about my user base, I don't want them knowing about my business model, I don't want my academic research to instantly commercialized, you might see the internet going dark in a very, very big way. And that might be a natural limiter, to the ability of these AI models to grow. And so the intent might be a very different place.

Erik: That's a really excellent point, because I stop and think about it, you know, all of the companies who are doing a good job today of putting all of their user manuals on the internet. So when you got your five year old appliance, you can go and find the manual that you lost six months after you bought the thing, and read up on how to change the time and date or whatever it is that you need to do. Well, if you think about it, a eyes are going to be able to scrape the entire internet, read every user manual for everything. And then project by looking at the the pace of how the features of five years ago as product compares to three years ago as product compares to last year's product. They'll be able to project what the features should be for next year's product, launch a company and deliver that product in order to compete with whoever they're studying, there'll be able to go in and that as you say, could create an environment where everybody's kind of on a need to know basis you you want to you want to use her manual for a Nikon camera, we'll prove to us that you actually own a Nikon camera before we give it to

you because we don't know you might be an AI and AI that's got some other motive. So that could really change the entire attitude on the net about access to information.

Matt: Exactly right. It could be a very, very different world out there in the next year or two in terms of an open Internet going to a dark internet.

Erik: Well, Matt, I can't say that you have completely allayed my concerns about the future. But I do want to thank you for a terrific interview. Obviously, this has been a long one. Before I let you go though, please tell us a little bit more about what you do at freelancer and particularly you've written a write up about your own research and experimentation with AI and so forth. That is linked for our listeners in the research roundup email. For people who don't have a research roundup email, just go to our homepage, macro voices.com. Click the red button above Matt's picture that says looking for the download. So you'll be able to get that download. Matt, tell them what's in that write up what they can find there. And also what freelancer.com is all about?

Matt: Well, I run [Freelancer](#), that's the world's largest freelancing and crowdsourcing marketplace. We have 67 million artificial intelligence sentient beings, being humans that can get any job done, you can possibly think of whether it's something simple like build me a website design from an app right through to gene editing or data science or crack propagation in satellites. We do all sorts of crazy things. So try us out if you want to get any sort of job done or you're building a business want to grow your business. And all those freelancers now are AI powered. So they're all using the AI tools to get the job done. So you can get amazing things in terms of the app of the worker at very, very low cost. But I think what readers will find very interesting is in collecting my thoughts for today's discussion, I put together a bit of an essay and it's in the research Roundup, it's a little bit long, but I think people will find it interesting if they want to kind of get up to speed very rapidly on where is the state of the art at least in now into June to July to August of 2023. In AI technology, where might it be going and what are the summit's some of the challenges and the risks and opportunities that might present themselves as a result of it?

Erik: Matt, I can't thank you enough for taking the extra time for this extended-length interview. I'm really looking forward to following your ongoing coverage of AI.

I'm going to wrap our summer special with my own prepared closing editorial monologue.

To be sure, the good side of AI is really good. In fact it's SO good that the biggest challenge will be figuring out how to absorb the societal impact of millions of jobs becoming unnecessary because they can easily be automated with robots and AI. So to be clear, the potential benefits and positive impacts of AI are hard to even fathom they're so great!

But I'm more focused on the various reasons I've become convinced that AI is now 2nd only to global nuclear war in terms of the existential risk it poses to society. I know many of you are rolling your eyes because you think I'm an eternal pessimist prone to always seeing the glass

half-full, and that's fine. I endured all hate mail from MacroVoices listeners who called me a reckless and irresponsible fear-monger for proclaiming on January 30, 2020 that a global pandemic was the most likely outcome, and I'm sure I'll get some more hate mail for these comments.

Experts who understand how AI really works would be quick to point out that we're still a long way from General AI, which is where the computer actually gains a sentience, or consciousness. Because that's still a long way off technologically, they are quick to dismiss laymen's fears about a Terminator-like scenario where the machines form an evil conspiracy to wipe out humanity.

What most people are completely missing is that you don't need sentience or a 'singularity', or machine consciousness in order for killer robots to destroy humanity. The generative AI technology we already have right now today is more than sufficient to compete with nuclear warheads to become the technology that ends humanity on planet Earth, and we don't need Terminators or Time Travel or even computers that are capable of sentience or consciousness in order to destroy humanity. We already have all the technology we need.

You might be thinking "Come on Erik, the most important people in this field are already raising red flags and they're going to be responsible and put limits and controls on AI so it doesn't get out of hand." I specifically predict that efforts to control or limit the power of AI will fail completely, and that people so inclined will be able to get their hands on unrestricted AI that has the features meant to stop it from being used to help terrorists disabled.

The reason I'm so convinced of that is that while all the business applications of AI are incredibly compelling, nothing comes remotely close to military applications in terms of AI being a perfect fit. Maybe too perfect a fit.

Any notion of limiting or controlling AI will be disregarded by military users, and for legitimate reason. The military rationale is that the cat is already out of the bag. AI exists, and the bad guys will eventually get their hands on it. So the question is we either use it ourselves to beat the bad guys, or we allow them to use it to beat us. If those are the choices, the "military mentality" of arming our military with a new weapon before our adversaries arm themselves with that weapon makes perfect sense.

At first it seems that any semblance of sanity would dictate never allowing the AI to make a decision to take a human life. For that reason, I predict that Version 1.0 of software controlling lethal military robots will be designed so that lethal force is not possible without human decision.

But that mindset won't last long. The reason centers on what military people call the *fog of war*. On a battlefield, your computers and radios and other equipment giving you important information get blown up from time to time, and bombs are going off left and right. Human beings are pushed to the absolute limits of their ability to cope with extreme stress on the battlefield, and many suffer PTSD for decades as a result.

A new technology that is immune to human emotion and which can literally see through the fog of war by performing millions of risk assessment computations could make a huge difference in saving our soldiers' lives, and the fact that more of our soldiers would die unnecessarily if they didn't empower the robots to make advantage of their ability to see through the fog of war. And that's how the policy of always requiring human direction before unleashing lethal force will eventually be abandoned.

Then we'll have super-fast, super-smart combat robots and drones that make their own combat decisions including lethal force decisions. It will work well and save lives at first, but sadly that will just embolden our military toward even more adventurism, since the cost of our own casualties is declining.

We don't need general AI or evil motives for AI to destroy humanity. All it will take is to first establish a precedent for allowing AI to make lethal force decisions in military robotics. Then someday someone will be doing a relatively mundane training exercise, in which the robot is programmed to try and get the most points by scoring the most kills in the war game. Then the computer, which has no sense of scale or consequence, will figure out that it gets points for kills, so the best way to get the most points is to have more enemies to kill. And the best way to do that is to take the information the AI read on Wikipedia about the Gulf of Tonkin incident, and use it for the AI to plan its own false flag to get the humans fighting with one another, so that it can eventually be unleashed to kill all the humans on both sides. Not because the AI got smart enough to have a consciousness and because it began plotting evil conspiracies, but simply because some guy programmed it to get as many points as possible, without stopping to consider just how resourceful the AI might be in finding ways to score more points than its designers ever imagined.

I also predict that it's only a matter of time before ChatGPT is cracked, meaning that hackers figure out a way to disable or reverse the programming it contains to prohibit helping terrorists and so forth. Once that happens, governments will legitimately become concerned that terrorists have a version of GPT4 that doesn't resist when they ask it to read every single word ever written about 9/11 and terror attacks generally, and then plan an optimal one for the latest Jihad. That concern will lead to much more government authoritarianism, restrictions on the Internet, and a general acceleration in the already-apparent trend toward greater authoritarianism of governments over the governed.

Bottom line, my conclusion is that the good side of AI is the best thing to happen EVER, but that it will be more than outweighed by the bad side of AI, which I really do think poses an existential risk to society on par with nuclear weapons.

One thing you can't afford to do is ignore AI. It's going to change all of our lives completely. For the better in some ways, and those should be the first we see, so let's enjoy them and hope I'm wrong about the bigger risks.

We'll be back to our regular show format next week, and that concludes this year's MacroVoices Summer Special. For the MacroVoices podcast network, I'm Erik Townsend.