

Understanding Artificial Intelligence: The Opportunities and Dangers of the AI Revolution



Jeff D. Opdyke

A Quarterly Wealth Advisory report

Understanding Artificial Intelligence: The Opportunities and Dangers of the AI Revolution

A *Quarterly Wealth Advisory* report

With Jeff D. Opdyke and Devansh

Cover photo: [istockphoto.com/da-kuk](https://www.istockphoto.com/da-kuk)

© Copyright 2023. Global Intelligence. All rights reserved. No part of this report may be reproduced by any means without the express written consent of the publisher. The information contained herein is obtained from sources believed to be reliable, but its accuracy cannot be guaranteed. Registered in Ireland No. 285214.

This report presents information and research believed to be reliable, but its accuracy cannot be guaranteed. There may be dangers associated with international travel and investment, and readers should investigate any opportunity fully before committing to it. Nothing in this report should be considered personalized advice, and no communication by our employees to you should be deemed as personalized financial or investment advice, or personalized advice of any kind. We allow the editors of our publications to recommend securities that they own themselves. However, our policy prohibits editors from exiting a personal trade while the recommendation to subscribers is open. In no circumstance may an editor sell a security before our subscribers have a fair opportunity to exit. The length of time an editor must wait after subscribers have been advised to exit a play depends on the type of publication. All other employees and agents must wait 24 hours after on-line publication prior to following an initial recommendation. Any investments recommended in this report should be made only after consulting with your investment adviser and only after reviewing the prospectus or financial statements of the company.

Introduction

On Nov. 30 last year, artificial intelligence research company OpenAI released an early demo of [ChatGPT](#), an AI chatbot that can generate a clear, cogent answer to almost any question it's asked. Within five days, it had attracted over 1 million users.

Students started using it to write high-school and college essays... programmers began using it to create workable code for websites... investors were using it to produce profitable currency-trading programs...

This marked a technological milestone as big as the emergence of the internet. It signifies that the AI age has begun.

Today, ChatGPT is attracting nearly 2 billion website visits per month. Meanwhile, the tech sector has seen a massive influx of investment as a result of this newfound interest in AI, with the Nasdaq jumping 31.7% in the first half of 2023, marking its best first-half performance since 1983.

The question now is, what comes next?

ChatGPT is an AI language model that generates answers using algorithms that have been trained on vast amounts of written data, much of it taken from the internet. But this is just an early example of the types of AI models that are in development...

Future AI systems threaten to fundamentally transform countless fields such as medical research, supply-chain management, data analysis, education... the list goes on and on. And as with the Industrial Revolution in the mid-18th century, this could upend hundreds of professions and displace countless workers.

Nevertheless, this tidal wave of innovation and disruption cannot be stopped. The AI revolution will produce new winners and losers on a massive scale, which is why it's vital to understand what this future looks like.

So, in this edition of your *Quarterly Wealth Advisory*, I interviewed Devansh, a software engineer and artificial intelligence specialist. He currently works as a machine learning engineer at tech firm SVAM International and writes a popular blog on the website Substack called [Artificial Intelligence Made Simple](#), where he discusses the implications of AI for our future.

In the edited transcript of our interview below, we discuss his views on the ChatGPT revolution... what the AI future will look like... whether the current AI investment trend is a bubble... and why the Google or Apple of AI likely exists out there right now.

Check out the transcript of our conversation below...

Sincerely,



Jeff D. Opdyke

Understanding Artificial Intelligence: The Opportunities and Dangers of the AI Revolution

Jeff:

We've been hearing reports about AI for the longest time... easily more than a decade. Then, suddenly, ChatGPT pops up and AI is everywhere.

It's the talk of the investment world. Everybody's discussing AI and ChatGPT and all the different ways they could be used in marketing, writing, graphic design, etc. What's leading to this? Why is it that suddenly we've gone from AI is in our future, to AI is here *now*?

Devansh:

To look into that, we first need to understand the context of how AI and machine learning—or the subset of AI that ChatGPT is a part of—have been progressing over the last few years.

[Editor's Note: Machine learning is a field of artificial intelligence that allows systems to learn and improve from experience without being explicitly programmed.]

As you mentioned, you've probably been hearing about AI for at least a decade now. It was 10 years ago that the Google-owned company DeepMind programmed an AI to beat human players in the boardgame Go.

Since then, we've seen a lot of investment come into the field, particularly in areas where knowledge is really tricky to decipher, like protein manipulation or supply-chain analysis, which I've worked on. Often, in fields like these, it was very hard to come up with an exact formula for how relationships work between the data and the value we want to predict. But machine learning makes that easy. So, we're starting to see more and more cases where machine learning is being used fairly extensively.

That forms the backdrop of the ChatGPT revolution, if you would call it that.

The reality is that we've actually been using AI models for some time that do very specialized things. We've trained them on very specialized data. For instance, if I want a model to predict the sentiment of a movie review, I would train it on the data of movie reviews. If I wanted a model to help me generate some code in the Python programming language, I would train it specifically on that.

But as we start to get access to more and more data, and as the internet becomes more filled with content, more people see it as a smart decision to turn to AI for all kinds of things. So instead of just making an AI Python generator, what if it did Python, Java, Go, and a bunch of different languages? In 2020, we saw this with Meta [Facebook's parent company], where they had some great results in translating languages between each other.

As these data sets and modules scaled up, the original ChatGPT came along. It could take a large set of the internet's data, encompassing a bunch of different things, and combine a whole host of different models. If a simple, smaller model can learn one task really well, there's no reason why a bigger model—trained with much more data on multiple different tasks—can't learn all of those tasks well. That's the rationale behind GPT.

However, when we look into this in practice, it falls apart a little bit. ChatGPT is not that great when you actually start to look into a lot of the systems and studies about it. The juice is often not worth the squeeze. But for the internet, it's good enough to go viral in a one-second click or a ten-second video. The fact that it can seemingly do a lot of different tasks that you wouldn't think possible... it seems magical. And that's how we got here, where we're starting to see people trying to use GPT everywhere, even on tasks it's not able to perform.

One influencer was saying GPT would kill search [engines like Google], which is actually a completely different task. Or that GPT could be used for text summarization. By the way, don't use it for that.

It's been demonstrated multiple times that even with simple text, GPT can hallucinate [deliver false answers]. We saw a lawyer try to use GPT to research a case and he came up with false information regarding that.

So, in all these cases, what we see is ChatGPT looks like it can do everything. And then most people accept that and they don't look into it. And that gives GPT a kind of mystique.

If you see everybody talking about GPT doing this, or GPT doing that, you get a bit of a FOMO cycle going on. [*Editor's Note: FOMO is short for "fear of missing out."*] Investors start clamoring for it.

I was with a startup and our investors were constantly asking us to use GPT even when it was not appropriate. But then you see companies use it, furthering the cycle.

Jeff:

So, you're saying there's a lot of hype and heat around AI right now. And maybe some of that hype and heat is not warranted simply because it's built around a bunch of influencers and media coverage basically giving artificial intelligence a little more credence than it's probably due at the moment.

That said, clearly AI is going to be a big part of our future from here on out. What do you see that AI future looking like? How widespread is AI's reach going to be when it's fully developed?

Devansh:

AI will be everywhere. And there's really no doubt about it, because pretty much everything from a simple "if" statement in a computer program that checks a few conditions, to these very complicated models, is considered AI.

One of the things that's led to this current boom in AI and machine learning is the fact that people have started waking up more and more to the importance of data sanitation, data processing, documentation, etc. [*Editor's Note: Data sanitization is the process of irreversibly destroying data. This process helps AI systems since it means older data is deleted rather than retained, so AI models have newer, more accurate data sets to draw from.*]

That's really enabling more data-driven decision-making, which enables more people to really step in and start adding value to the data, because now we have enough comprehensive, high-quality data that we can work with. And that's really where I see AI being a great benefit to society.

Traditionally, a lot of the processes and decision-making in our societies have been somewhat arbitrary, but now we can now leverage processed data to actually start making informed decisions. And that's really where AI would be extremely beneficial... augmenting human data and human decision-making capabilities, because it can extract so much from a lot of information.

Jeff:

What do you see as the downside of AI in the future?

Devansh:

The downside—as we've already started to see, and will continue to see—is that these fields are so complex that there is often a large disconnect... even among people working in these areas.

To give you a very simple example: If you were to bring me in to evaluate employee productivity or candidates, I might be able to do all kinds of statistical analysis to help you understand trends that would be helpful. But often what we see is there's a large disconnect between the people who do the statistical analysis and the people who understand the core domain. And that can lead to situations where we perpetuate biases that we really want to avoid.

So, with this example—and I've seen this a lot recently—let's say we're going to use GPT for hiring: resume and candidate evaluations, etc. Or we're going to use GPT for evaluating if our candidates are doing well, if they're engaged, etc. The problem with that is GPT, by the nature of being trained on the internet, encodes a certain kind of voice. It's a very strong American, white, affluent, neurotypical voice. And I'm not saying that to be “clickbaity.”

That just happens to be a large part of the internet, which it's been trained on. And in cases where we had to evaluate candidates, what that means is the system would unfairly penalize candidates that are unable to play to these stereotypes.

So, the problem is, when you bring in an AI analyst like me, I wouldn't have the context to understand, “Oh, look, this is not good. This [system] will have these flaws.” I would not even think of how it would evaluate these neuro atypical people or neurodivergent people. I would just do an analysis on the data, and I'd say, “Hey, according to the analysis, these traits are problematic.”

Then the data-driven decision-making would start to further isolate these groups. And this is the real problem with AI and AI safety and security.

A lot of the devils of these problems are overlooked because the data guys like me will focus on the technical challenges of studying all this data. I will wonder whether to process it this way or that way. But I won't necessarily know the context behind the data.

Then somebody else, a business user, might just look at the end results and say, “Oh, this person was able to identify weak performers much better.” But what it doesn't identify is that the system that's evaluating these weak performers is alienating certain groups, and that's going to be worse for social inclusivity in the long run.

Jeff:

How do you fix that? How do you train AI so that it isn't just sort of a white, Eurocentric version of the world, for lack of a better term? How do you make it a more globalized system?

Devansh:

The problem is not that GPT is white, Eurocentric, etc. The problem is that we're using it in ways that would be problematic. GPT can be beneficial in that it can act as a mirror for the internet... to see what kind of biases there are in the training data set.

For example, when it originally came out, to prove a point, I asked ChatGPT to rank the smartest racial and gender groups. And it didn't mention Indian people at all. It did mention Latinos, but not Indian, which is a very clear indication that a lot of the data has a very heavy American bias to it.

Because if you looked at it just from a pure statistics perspective, there's a lot of Southeast Asian and Indian people walking around the planet. If this was a purely representative sample, we would not see Indian people excluded.

So, the problem isn't that GPT is biased. The problem is that it has very surface level capabilities. It can write a cover letter, or it can tell you this candidate uses a lot of conversation fillers or whatever, implying their communication isn't great. But it doesn't give you the context behind that.

The real problem, then, is that we as people and decision-makers need to step back and really identify whether certain use-cases for AI are appropriate in the first place. And that's been the case with AI and machine learning for a while now.

We need more voices so that we're aware of these situations. I'm only aware of this particular kind of use-case problem because I've written about it a little bit. So, I've had people reach out and tell me, "I feel very excluded based on the metrics."

But that is really where it needs to go. We need to really look and evaluate human systems at a deeper, fundamental level, as opposed to directly jumping ahead and trying to integrate AI systems.

Jeff:

When you're looking out into the next decade, as AI becomes a bigger part of society and business, what industries do you think are going to be most transformed by AI?

Devansh:

Generally speaking, medicine, healthcare, and medical research are really going to benefit.

When I lived in Memphis, I lived with a researcher who was studying alcoholism. She used to get rats drunk, basically, and see how they reacted to getting drunk. And I sat in on one or two meetings out of curiosity, and then I spoke to her research supervisor. They were using what to me was very, very rudimentary machine learning and AI.

I asked him, "Why don't you use more advanced systems?" He brought up a really valid point to me, which was that in their field, they were the first generation of people using any kind of machine learning.

So, they had to set up the data sets and everything else from scratch. This is what's currently happening. Most veteran labs have their first generation of data analysts and data scientists. The researchers are still setting up their data experiments, setting up the validation and so forth. As this becomes more mainstream, as these processes become more refined, they will start to have access to high-quality data sets on which you can do the real value-added analysis. And that's when we'll see the field explode with growth.

Any field where traditionally we haven't had very great data handling or data processing, AI investment will come in. The one good thing with the GPT hype has been that people are now paying attention. I spoke to a sociologist recently who was looking to integrate more data sanitation and data cleaning into her processes. So, all of these use-cases will blow up. The social sciences, medicine... all these fields where traditionally data science didn't play as big a role.

In fields like finance, we have always associated them with data. So, AI will have a lesser impact on those, in my opinion. But fields that haven't yet been revolutionized by data will really blow up thanks to AI.

Jeff:

Are there industries that probably won't be transformed by AI?

Devansh:

That's a good question. Off the top of my head, fields that probably won't be transformed, if I had to pick some, would be anything that has a very strong human empathy component to it.

If you were taking care of old people, I don't think it's an information issue as much as it is an issue of being able to deal with people, being able to manage them, calm them down, treat them properly. And sure, I can see things like Fitbits or devices to monitor health better being able to identify issues. But I don't see how its core responsibility, which is taking care of other people, could be automated away easily.

Jeff:

In that context, what if you had some sort of robotics combined with AI? And you have a humanoid robot that is programmed with AI, that knows how to deal with the elderly?

Or what if you had some sort of robotics that's AI-powered that knows how to manage a customer at Burger King or McDonald's? And you have robots in the back that are making fresh burgers and things like that? I mean, it seems to me that if you combine AI and robotics, you negate the need for most workers in the world.

Devansh:

Sure, but that's also very expensive and very hard to do. We've been talking about AI taking away the jobs of burger flippers for longer than I've been alive. I definitely know that conversation started in the '80s, perhaps earlier.

What we see is, it's not these jobs—like, say, electricians, plumbers, all of these jobs whose tasks are so fine-tuned and require motor skills—that are going to be replaced.

AI is trained in a 2D world. It's trained on a computer and essentially all the information it gets is in the form of a pixel, which is then transferred over to a 3D environment. That process is a lot harder than people think.

In 2014, we had very prominent people like [computer scientist] Geoffrey Hinton claiming that AI would replace radiologists because AI can seemingly predict something in your brain better than radiologists from looking at images. But that never came to fruition. Not because the radiologist lobby is very strong and they've been suppressing the plucky little AI, but simply because when we start putting things into practice, we encounter technical challenges that very rarely can be solved by AI.

Often, it's going to be things like... if I'm a burger-flipping robot, the moisture is going to really get out of hand.

Or if there's burger-flipping robots, what's stopping people from throwing things at them and hopping over the counters or doing any of that other stuff? And then as you start making the stores more secure, the costs would get so out of hand. The wages you'd pay your workers would be less than the cost to replace them.

So, there's also the financial aspect to consider. Theoretically, with enough money and resources, I can see all these jobs being replaced, but is there an economic incentive to do so? Because ultimately, how much are these [AI robotics companies] really going to charge? If one of those robots breaks down, that's \$100,000 gone. It doesn't cost \$100,000 to employ a burger flipper. We also have to consider it from that perspective.

Jeff:

So that segues into one of the things I'm curious about: the investment angle. Everybody's rushing into AI-centered stocks these days, like Nvidia or Microsoft. In some ways, this hype reminds me a lot of the original internet boom in the late '90s, which I was part of when I was writing about it for *The Wall Street Journal*.

Back then, there were all these new metrics out there. People were valuing companies based on how many eyeballs they had every week on their website and weird things like that. And a lot of those companies ultimately failed.

I get the feeling we are in that phase right now with the AI hype... that there's a lot of companies putting AI in their name or they're becoming some sort of artificial intelligence company, but at the end of the day, they're probably just more sizzle than steak.

So, I'm curious if A) you see it that way and B) does the "Google of AI" or the "Apple of AI" or pick one... do those companies exist yet, or is it still years before we find the true unicorn winners in this space? [*Editor's Note: A unicorn is a startup that reaches a \$1 billion valuation.*]

Devansh:

"More sizzle than steak." I like that. Absolutely. That's very much how it is. When markets see something new, we see a lot of money being thrown into ideas that are fundamentally not useful.

A bit of a hot take, but Uber is the perfect example of this. Investors saw the profitability of platform-based systems like Facebook and then they put so much money into other use-cases [like Uber or Doordash], which they thought would be platform-based. But we found out that their economic structures work differently.

With AI, there's even more hype. Most of these companies don't have a differentiation factor. Most of these are just front ends on a GPT that raise some money... that are trying to gain some attention. And I've spoken to quite a few.

I think the next big AI player or winner is probably already out there, and we're going to see them come forward eventually, because I think the next big AI player will be the one that enables Level 1 AI for organizations that want to get it. [*Editor's Note: Level 1 AIs are basic AI systems that can automate simple, repetitive tasks.*]

So that's what I see as a big opportunity: startups or businesses that basically help other teams set up their AI and monitoring platforms. And those platforms won't be the cream of the crop because they are general purpose setups, but they will really help organizations come into this space.

For example, the group I'm currently working with, SVAM International... what we do a lot of is helping organizations set up their processes, set up their pipelines, monitoring, all of that. And there are other companies like Lightning that are very focused towards setting up AI pipelines for different organizations.

These could be the real big winners. I don't know if you've read the book *Jugaad Innovation*, but one of the very interesting business lessons in it was to integrate the margins. Instead of chasing what's already happening, try to target the people who are not traditionally part of the system because that's an area you can monetize... kind of like a blue ocean for you. And that's really what I see being one of the big AI opportunities: helping other people set that up.

But again, that becomes a complicated process because there's so many people who are offering to do just that. So, quality can differ substantially.

Jeff:

Have you invested personally in any AI stocks?

Devansh:

Just the standard ones. For me, personally, I treat my writing and my work as a business or startup, per se. So, the way I see it, most of my time and energy go into that because I would rather make that work. My AI investing right now is fairly rudimentary, fairly basic, though I used to do that earlier.

I identified stocks like Cerence (CRNC) prior to 2020, where I saw some of their research and was like, "Hey, the computer vision looks really cool. I think this is going to be big." So, I invested in it. And there was around a 4X or 5X return on it.

There were a few others. But as of this moment, I don't see any others like this.

Jeff:

Final question I want to ask you: at some point, doesn't AI kill AI jobs? I mean, we have all these people who are pitching AI marketing services and AI writing services and AI this and AI that.

But at some point, it would seem to me that AI gets to a point where some developer is going to write the code that basically teaches AI to *teach itself* how to do all these things. So, at some point, doesn't AI kill AI jobs?

Devansh:

Very unlikely. Because again, these AI solutions only give you the illusion of competence and the illusion of what you would pay somebody for.

Let's look at it from a professional context. ChatGPT might be able to code better than my mother. But you would not have paid my mom to code in the first place. So that's really what we see with a lot of these situations. Sure, even if it writes better than the average college student, the average college student wouldn't get paid to write about a specific topic.

We've already started to see this massive burnout and backlash in the content community. You can very easily start to identify things that are AI-generated because, again, they don't have a voice of their own. By their very nature, they have encoded biases because they are basically averages [of the existing data].

What you've done is you've created the most average writer possible... the most average coder possible. And they're going to fail with simple tasks. We saw this with ChatGPT. Best of luck trying to explain to GPT the requirements of a complex system.

What AI will do really, really well are the things that humans typically aren't good at, which is retaining a large amount of information, or being able to go through something like a large legal document with a fine-tooth comb to extract information and pull insights. It's going to do that phenomenally.

And I'm sure there are jobs that are largely that, which will be either replaced or be changed so one person would do the job of 10.

But the higher levels of analysis, the higher levels of synthesizing multiple sources and putting them out together, that's an area where we aren't even sure how we'd build an AI to do that because we can't even train AI to do the simple things reliably, in many cases.

You can't just keep scaling AI models up, because even as you scale these up, there are very clear experiments that I can do to prove this AI doesn't understand what it's doing. This AI doesn't really understand the context behind it. It's essentially just statistically associating things together and churning them out.

So, I wouldn't be too worried about AI killing AI developers' jobs. I would be somewhat worried about people *thinking* that AI is going to kill AI developers' jobs.

Throwing AI into a use-case like we saw with JPMorgan and the WADU [Workplace Activity Data Utility, an AI employee monitoring system introduced by the bank] made life meaningfully worse for everybody. That's a much bigger problem.