VinciWorks

AI and Business

What companies should be thinking about when they think about AI

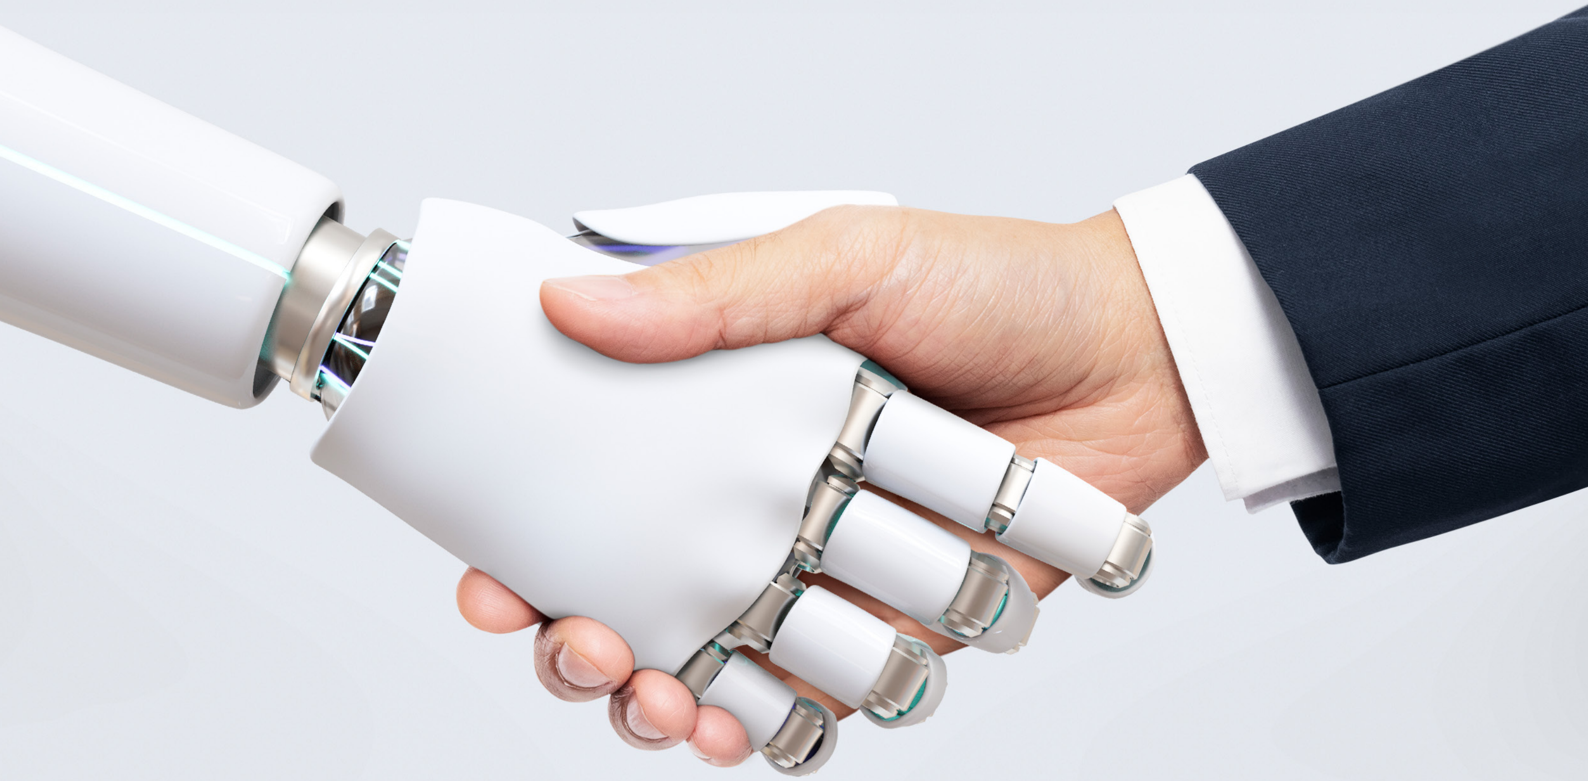

A conversation with Shlomo Agishtien,
AI lead at Trullion

Introduction

As artificial intelligence (AI) tools are increasingly becoming part of the daily processes of nearly every company and AI regulations are bearing down with the EU's AI Act leading the charge, it's more and more important to understand how to utilise and develop these tools ethically and effectively. Naomi Grossman, learning and content manager at VinciWorks, sat down with AI expert Shlomo Agishtein to discuss what companies need to understand about AI, how these tools can be used, why an AI company policy matters and how worried we should all be about AI regulation.

Shlomo Agishtein is the AI lead at Trullion, an AI-powered accounting solution company. Three years ago, before artificial intelligence (AI) was the big topic of conversation, Trullion had a vision to bring AI to accounting and auditing to help automate these processes. Shlomo has led that charge and in the course of his work, he uses many different aspects of AI, including different areas of machine learning (ML) and AI data science to be able to bring effective solutions to his clients. We sat down with Shlomo to learn how AI is changing the business world.

Naomi:

More and more people are using AI in their daily lives, whether they realise it or not. How is AI being used?

Shlomo:

AI is incredibly ubiquitous and, I think, very misunderstood. People think that AI started with Chat GPT, Gemini and Bard. But that's not correct. Google photos have been using AI for many years. So does Nest that controls your thermostat. And Waze. In fact, almost every single application you have uses data driven machine learning technology in order to be able to give you functionality. Think about Amazon Alexa or Siri. These are all using AI.

Naomi:

Is it important for people to actually understand what artificial intelligence is?

Shlomo:

I think that there is a lot of value to it. What it means is understanding that, for example, ChatGPT is not actually searching the internet or understanding that when companies train models on data, it's not that there's some database that contains your data and it could be leaked. If the data is trained on your “model,” it means that the model is made of weights (which are numbers) that are tweaked to make the model's predictions better. So there is no 'data' actually stored in the model; it is just used to move those numerical values in the directions that make them better predictors. There's absolutely no way for someone to get your data from the model weights. Also, large language models (LLMs) have been used to generate data that they have seen in training, but for the vast majority of models this is not relevant.

Being educated on the differences between deep learning and explainable AI and understanding what these different terms mean and how they're being used can make you an educated consumer and an educated user and somebody who is able to really manage this new world.

Naomi:

How significant a role do you think AI tools play in business growth?

Shlomo:

And that’s not in the long term, that’s in the short term. Obviously we are an AI company but customers who are using our products are saving days of time when they are using AI to automate tasks. It used to take auditors 30 to 40 hours and

> ”
>
> **Any business that's not trying to formulate some sort of AI adoption or AI strategy is going to suffer.**

now it takes them 10 minutes. But even on a much smaller scale, if you're doing research or you're trying to generate content or you have these large databases which store data that nobody knows what's there - there are so many areas where AI can supercharge your efficiency. It's really important to educate yourself about what's out there and formulate some sort of AI adoption strategy. It doesn't just mean you adopt it all right away. There are all sorts of check boxes and protocols that you need to consider very carefully before you decide on what to use and how to apply it.

Naomi:

Can you give some specific examples of what you've seen?

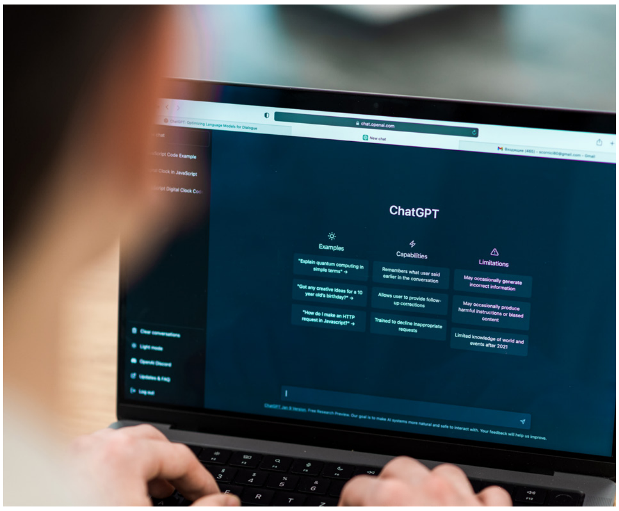

Shlomo:

There are companies that are really disrupting every single area of existence from finance to marketing to healthcare. And they're doing it right now with very low hanging fruit of AI. I'll give just a very simple example from my work. One of the more difficult questions that auditors face is they upload a large amount of audit evidence and then they don't really know what they have there. But now, with not even a lot of deep software development, you're able to create a tool that can ask questions on the data that they upload and get natural language answers which will also retrieve the relevant documents with the data that they're trying to find. Another example is search, which was always available, but now it's able to understand. It's not just looking for keywords or similar words but it can have a deep sense of the meaning of the text and what you mean by your question, and it can match those documents that are relevant to the meaning of your question. That is something that is very powerful and very new and it's a massive efficiency boost.

Naomi:

What are the top AI tools that companies should know about?

Shlomo:

It really depends on every company's use case and industry. There are companies now evolving in every industry to address common pains using AI. It's rare that you're going to have a general out-of-the-box tool - and those also tend to be risky because you don't know what's happening to your data. In general, ChatGPT4 is very powerful. If you just have a subscription, it will make a lot of aspects of your work easier. You're drafting an email you don't know how to write or you're a programmer and you have a function that's not behaving, you can put that code into ChatGPT4 and it will help you find bugs. You're a lawyer trying to analyse a contract and there's some passage that you're not sure about, it can help. You can also bounce ideas off it. Of course, there are hallucinations and all sorts of risk factors to be aware of. But programmers are known for having rubber ducks on their desk for whenever they want to be able to talk through a problem. This is the rubber duck except it really interacts with you. So the real point is understanding your pain and then understanding what companies are offering.

Naomi:

There was a tremendous amount of interest and panic around GDPR. Do you see GDPR impacting AI growth in companies?

Shlomo:

I think the biggest challenge of GDPR for AI development growth is this notion that people have a right to know how their data is being used. And I think that's a really good thing. People should know how their data is being used, but explainability in AI is a very big challenge. And this combined with the AI Act is pushing towards really understanding and being able to explain the AI process.

The challenge is that a lot of the really advanced AI is kind of a black box and if you try to use other AI systems to explain these AI systems, then you end up getting yourself into all sorts of interesting situations. So, in a way it is good because it is pushing for a more explainable, predictive, deterministic AI that is able to walk you through a start process and at least give a sense of "okay, we rejected this because it features this but if you change your AI system in this way, you can get accepted." So, is it holding back AI progress to some extent? Yes. But that's not necessarily a bad thing.

Naomi:

We saw GDPR set a standard in data privacy regulations. Do you think that the AI Act will have a similar impact, or maybe even a bigger one, on enforcement? And will that have an impact on innovation in AI?

Shlomo:

The AI Act is very interesting in that it took a risk-based approach to regulating AI. I think there was something clever in that because it helps set the stage for the fact that we can't predict AI development but we can predict harm. We can think about whether it's monitoring people's social behaviour or if it's doing mass facial recognition. That is a harm that we are able to identify a lot easier than technologies. I think that the AI Act in that sense is something which is a good step forward instead of trying to regulate tech. Regulation obviously is necessary. But I think that the biggest fear is who it helps and ideally, we want regulations to help people. But if you're going to shut down open source, if you're going to shut down the ability of startups to be able to develop products quickly because of the regulatory burden based on them, you're really just handling the keys to the kingdom to companies like Google, open AI, Microsoft and Facebook, who have the ability to be able to deal with these regulations, with the lobbyist, with the lawyers, with everything they need.

And in some sense, then you're also at the point where regulations could end up becoming almost useless because these people know how to deal with these types of regulations and still do whatever they want. That's why I think, there has been efforts by people that were working with AI to really try to make it open source friendly because open source is very important. You can see the code. It kind of breaks the control of this kind of monopoly that is really not in the best interest of people. So I think it will cause hurdles for people especially in industries like ours where we have to think about financial data and regulations. But I think ultimately [the AI Act] is a good first step. It's really just a question of how nimble the EU is going to be with adopting it as the various challenges to it arise.

Naomi:

What do you think about Sam Altman from OpenAI asking the US Congress to regulate this technology?

Shlomo:

And I think that if the government wanted to really ensure the future safety of AI, it should be more of a proactive rather than a defensive approach. And we should be thinking about it from the terms of AI as the new infrastructure. We should be thinking about it like we thought about the Hadron collider or the internet or all these things where there should be massive government investment in building really powerful AI models that are monitored by the government. And, then people can buy cheap licences for them. This is how you can allow private development. The existence of strong competitive alternatives that are safe and monitored and are their primary accountability is in the public interest. I think it will do a lot more than trying to figure it all out and trying to police a technology that nobody in Congress really understands.

> **The problem with the Sam Altmans and the Googles of the world going in there and saying AI could be super dangerous and we have to regulate it is that, in some sense, it's self serving because if you regulate AI, you are limiting who could build it.**

Naomi:

How can companies effectively manage their data privacy issues when they are using third party AI providers, users or developers?

Shlomo:

That's a very sensitive question and this is going to be a controversial answer. They don't. Obviously, you should make sure that the people you're dealing with have the necessary compliance and that they're compliant with data privacy and data security. If you're using the cloud you need to know where your data is going and what it's being used for. You should educate yourself about the data. But I've heard from companies, especially more traditional, risk averse companies, that are really scared to allow their data to be used to improve AI models. And the reasoning for that is security. They're worried about leakage or other problems which I'll talk about shortly. And they also think of their data as their way of protecting themselves, that they have this data, it is valuable and it gives them a competitive advantage.

But as far as safety goes, the data is going through the systems anyway. And as I noted, the fact is that it can be used to improve processes. There have been bad stories where LMS have spun out proprietary information and that's valid. But, in general, for most AI applications, it's going to improve the results for you. The AI development cycle right now is extremely fast and it does not require the massive amounts of data that it used to require. Let’s say you're a retail company and you're working with an AI company that's trying to create tools that work well for retail. If they're a startup, they don't have real world data. If you give them your data, it could really help them iterate quickly and improve products for you. But you might say I'm only going to allow these products if they work for me. In this way, everybody is trying to be kind of selfish for themselves but they can end up hurting themselves.

Data is very valuable for AI and real world data is extremely valuable. There will come a time when we won't need a lot of data at all and everything can be hyper personalised and your data can only be used for your model. But in general, be very careful about your data, be educated about what happens with your data and also understand that if you want to contribute to you having better tools, we do require your data. This is not necessarily going to always be true. We're seeing that generated data is being used more and more to train AI models. But real high quality data is always going to be important and it only benefits the company.

Naomi:

We know that AI tools could possibly lead to discrimination within companies.This could be because a company doesn’t have clean data or their historical data had some biases. Do you think that that's actually an issue? And what can companies do about it?

Shlomo:

Understanding that your model is only as good as your data is an essential axiom in AI. If you're going to have data that's been sexist or skewed there is absolutely no question that your models will become sexist. There are two approaches to deal with this. The first is that you need to use models to handle tasks that can scale. For example, how do we do loan approval? One person used to have to sit down and go through all the features, all the data, right? We know how to do that. But the problem is when we're trying to approve a million loans, then we use AI because we just can't scale. But what you need are people to create what are called validation data sets. You have people who take a real random sample and there are plenty of statistical methods, and then you collect good clean data on what the distributions should look like, like how many women or people of colour or people from various socioeconomic backgrounds. And if you see that your model within is deviating too far from that percentage, then you have a

problem. That's kind of a post-fact check where you say, I have identified that this model is showing a bias because through real world data analysis, I see the deviation and then you have to go back and you have to back test your data and see if the data is biassed.

That is where I'm very hopeful about the power of generated data. If you can find what you think is an accurate data set that's small, and you are able to train that model to be able to generate more data based on that data, then we can know how to perform and we can do constant updates and perform checks. It's much more secure and you have more control over historical devices, random noise and other things in your data. But if you don't have that option and all data is suspect, then you need to be able to have rigorous tests that are able to check the quality of your data, look at the distribution, see if it's what you expect, see if there's drift over time, and consider how far back you should you be looking. You have to have a robust or at least a decent data science team that's able to grapple with that.

Naomi:

Do companies using AI tools need to protect themselves more in terms of cyber security issues? How can they?"

Shlomo:

It's always the battle between the sword and the shield. On one hand, there's been a proliferation of more and more sophisticated cyber security issues because fishing is now becoming a lot better. People could really be interacting with someone they think is their friend. Being able to write more and more complex viruses is becoming easier for people with fewer and fewer skills and there are all sorts of social hacking that are becoming much easier. On the other hand,pattern recognition is getting better so we are able to identify threats and to process a far larger amount of data to look for anomalies. And for most of the cybersecurity threats that are on a software/ hardware level, you just have to buy the right tools. So social hacking is a much bigger threat than people reaching out and asking for your password or asking you for private information. People are now giving a deep fake of your boss giving you a phone call and asking you for the emergency login. There could be some software solutions for that, but not much. It’s really about educating yourself that the person on the phone that sounds like your boss or on the video that looks like your boss may not be your boss. And that is a really terrifying realm.

Naomi:

Do companies need to train their employees on using AI tools? Is there any kind of danger zone they need to be aware of?

Shlomo:

There is no question that they have to, especially, if their employees are using open source or other AI tools. There have been stories with Samsung, with JP Morgan and other companies where they banned the use of generative AI models because private information got out. There are companies that don't allow you to use what's called Copilot, which is GitHub's LLM, which generates code. I was once writing code and I was trying to write a file path and in Copilot I generated a full file path from somebody named John to his dissertation research. I have no way of tracking down John and it's not super valuable information but it does generate code. There's a lot of conversation now about security and what data to put in, what data not to put in. Users have to be very careful to abide by company AI policy in general because there are two versions of open AI that are relevant. There's what's called the API, which is how programmers interact with it. And there's one set of protections there that open AI says, they're not gonna be using your data. They're not going to retain your data, there, all sorts of protection. Then there is the chat interface that you go through and there, it's not the same thing. It's a customer. It's a free application and they're using your data to train. It's important to understand that just because your company may be using an open API, it is not the same thing as you using your own private chat and putting in proprietary information or other kinds of data that is really not okay for it to get out."

Naomi:

Let’s talk about a company's AI policy. Who's setting up these policies and how do companies know how to set these up?

Shlomo:

I think it needs to be a real partnership between a technical person and a legal person. Ultimately, they need to know that they don't know and they need to research. I've seen cases where companies have come up with AI protocols that are very random and that are in the interest of safety, but they're just completely cutting off their access to innovation. They really need a dedicated team that's made up of technical people who can understand the AI tools and the data flow and legal people who understand the various ramifications. And I would say business people who understand what risks are going to hurt the business and what AI risks are not. It’s not just, if I use it, I may have issues. It's also, if I don't use it, I will also have issues because I will be behind everybody else in my industry who probably are going to use it. It's really important to have representatives from those three teams that are dedicated to this and really working to educate themselves and come up with a policy that is extremely adaptable, meaning nothing should be set in stone because technology changes, and, as you learn new things, business changes, competition changes and you'll see the standards in your industry will shift. First movers will get rewarded and get punished. So there's really a lot of flexibility that has to be built into this process.

Naomi:

What does the future hold for companies in AI? Where is this all going?

> **Anybody who tries to say they know where this is all going is lying. AI is going to change the world and change itself in a way that we just don't know what's going to be possible in a year.**

Shlomo:

But I will say that we are at a place where AI research and AI development is more and more into the world of autonomous agents. That means that entire units of work are going to now be able to be carried out by AI agents, which means entities that are able to function with a certain level of autonomy and they're able to make decisions and run processes and do things using AI. It's already here but the reliability and the sophistication is not there yet.

There's also an entirely different part of AI called reinforcement learning that's been developing over the past 20 years. It's been waiting in the wings for really robust agents to be able to be fully developed because it works with how these agents are able to function in an environment. It's the area of AI that deals with autonomous vehicles and game playing, where agents are placed into an environment with a set of actions, and they work to be able to choose the best actions based on their environment. But everything depends on the sophistication of the agent and now, with LLMS and with these really robust agents, the field of reinforcement learning could do really cool things. What that means practically for a business is that more and more processes, and not just data searching or checking, but actual procedures, are going to become automated - things like running an audit or handling certain types of software development tasks. There will be more of a shift into strategic thinking, big picture development and compliance. There is going to be a disruption. People who worked in procedures will obviously be impacted. But there will also be an incredible burst of innovation because now it will require less and less technical or industry specific skills to be able to achieve outcomes. There will be an explosion of startups and companies where all you need is an idea and then you can build it. There's already talk right now of the first one-person unicorn where a single person is able to build a billion dollar company. That's on the horizon because you have no-code tools where you can build applications. There's content generation for marketing with AI. There are customer success tools with AI. There are so many things that will enable us to be able to do things that we weren't able to do before.

Naomi:

It's a cool future. But is it also a little scary?

Shlomo:

There is no question that it's going to be a very different future. It's definitely scary because there's going to be a lot of dislocation. But if you think about the industrial revolution, and after that, there are always these massive changes. People have suffered from them at each stage. A normal person thinks about themselves and their family and doesn't always think about what's good for humanity in general, which is understandable. So individuals may suffer, people may lose jobs. There will be re-training, especially with people who are older, and there are areas where this could be very impactful but it depends on how it all plays out. It's unpredictable. I don’t think it's going to take over and we are all going to be living in poverty except for five people who will be wealthy. If anything, I think it will allow people to use these innovations to improve their circumstances in a way never before possible. But I don't want to be utopian about it either. There are a lot of risks, but I think that we're going to be very surprised in 10 years to see what the world looks like and who gained and who lost. And, there will be bad things that happen just like there were bad things that happened with the internet and there are going to be amazing things that happen like there were amazing things that came with the internet. We're going to have to manage the bad and appreciate the good.

Naomi:

That's a perfect ending. Thank you.

Shlomo:

Thank you for having me.

VinciWorks

Contact us

www.vinciworks.com
enquiries@vinciworks.com
+44 (0) 208 815 9308