



# Artificial Intelligence: Sifting the Facts

**Artificial intelligence continues to attract strong interest from investors and entrepreneurs alike, even during the global pandemic. But where is the sector headed? Which recent developments have been most overhyped or underplayed? In this AI Roundtable, *Coller Venture Review* speaks with IBM’s Aya Soffer and Wharton’s Kartik Hosanagar to do a reality check.**

**D**espite the global COVID-19 pandemic and economic crisis, artificial intelligence (AI) continues to attract strong interest from investors and entrepreneurs alike. According to CB Insights, a New York City firm that monitors startups and venture capital, while AI deals declined in the first quarter of 2020, funding “jumped by 51% from the previous quarter to hit \$8.4 billion.” Successful IPOs by AI-powered startups such as the insurance firm Lemonade—whose market cap soared to \$3 billion when it went public in July—have added more sizzle to the sector.

As this momentum continues, several questions arise about AI and where it is headed. Among them: Will AI be as transformational as, say, mobile or cloud computing? Which recent developments in AI have been most overhyped or underplayed? What challenges in AI deployment are unique to enterprises as compared with consumer applications?

*Coller Venture Review* discussed these questions and more at a recent AI Roundtable meeting with Dr. Aya Soffer, Vice President of AI Technology at Haifa Research Lab in Haifa, Israel, and Professor Kartik Hosanagar, John C. Hower Professor of Technology and Digital Business at the Wharton School of the University of Pennsylvania, who oversees the school’s AI for Business initiative. Dr. Leslie Broudo, Head of the Coller Institute of Venture at Tel Aviv University’s Coller School of Management, moderated the conversation. Following greetings and brief re-introductions, an edited version of the discussion appears below.

This article is associated with a companion piece by Dr. Lior Zalmanson and Professor Gal Oestreicher-Singer, both lecturers in Technology and Information Management at the Coller School of Management at Tel Aviv University. 📌



**Broudo —**

Will artificial intelligence be as transformative as mobile or cloud computing?

**Soffer —**

I think AI will be as transformational as mobile and cloud computing but maybe in a slightly different way. That is because AI touches everything. Part of the transformation in cloud and mobile—and enabling all that—is due to AI. AI is the underlying capability that has made all these things transformative and will continue to do so even more moving forward.

If you consider mobile, there are things we like and others that we may not like as much but are very



“  
**AI touches almost everything we do from the core of creating the technology all the way to user interfaces**

”  
Aya Soffer

necessary. One reason that mobile has taken off more than we had envisioned is because we can now know people’s locations. That is AI in the mobile environment. That is why “mobile” technology can now provide users with an experience that is personal. Even though sometimes we are surprised (or possibly dismayed) at how much these technologies know about us, that is what makes them extremely useful. That is on the front end.

If you think about the back end, AI holds the key to a lot of what we see that enables the miniaturization of mobile devices and what we see on how the cloud runs. We use AI to make sure the cloud is running smoothly; we use AI to predict failures and fix them in advance. In my view, AI touches almost everything we do from the core of creating the technology all the way to user interfaces. In the future AI will be even more transformative when it will allow everyone—even those without computer skills—to interact with technology.

**Hosanagar —**

If we rewind back the last 20 to 25 years, we can look at technologies that have had a huge impact on business and society, such as the Internet, the Cloud, and Mobile computing. There also have been other technologies that may have received a lot of hype but which have failed to deliver. To me, AI belongs in the first bucket with the Internet and Mobile computing.

When we ask whether AI belongs in one bucket or the other, I think about it in a couple of ways. First, how relevant is the technology across multiple industries or modes of our life. On this count, AI is fundamental. We see AI applications in health care, finance, education, professional services, manufacturing, retail and so on. The scope is extremely broad. Second, over the last 20 years we have seen an explosion of data in all aspects of our lives. We could not have mined that data 20 or 25 years ago. Now, we not only have data generation happening at an amazingly fast clip, but we have also seen machine learning progress so much that we can analyze that data and make sense of it.

A lot of factors have come into play at the same time, ranging from data generation to data processing to progress in machine learning algorithms to the fact that this is happening across industries. If you look at Cloud

or Mobile or the Internet, they have been transformational in terms of touching many industries and many aspects of our personal lives. AI is similar. It is advancing so rapidly that there is little doubt in my mind that it is significantly transformational.

**Broudo —**

To summarize, from both your perspectives—the technology-enabling component and the change in industries it is affecting—you believe that AI is highly significant. When you read the popular press coverage of AI, what do you think is being overhyped? And what is being downplayed to the degree that we do not appreciate what is going on?

**Hosanagar —**

If we look at the popular press, I believe where there is lack of understanding is the view that AI is almost magical and sentient, and it can be viewed in the same way as human intelligence. There is also this concept of super intelligence. In reality, AI today is what we can refer to as weak AI or artificial narrow intelligence. That means we give a machine learning algorithm good data on one specific task and we can figure out the patterns that allow us to make predictions on that task. For example, we can give data on whether an email is spam or not, and it can do a great job of figuring that out. But that does not mean it can be truly intelligent and transform itself into a robot that starts moving about in the physical world. You might have an autonomous vehicle that drives around, but that does not mean it can do other things. That is where we are.

Sometimes, when you read articles in the popular media, you start to see the discussion around AI suggesting that we have created something that is truly intelligent and mimics human intelligence. We are nowhere near that. We might get there in the future but we are not there yet. Similarly, there is a misperception that AI can beat doctors or trump their medical knowledge. Here, too,

AI is good at narrow tasks such as reading X-rays and other such tasks in radiology. But to believe that AI is smarter than doctors makes it seem like general intelligence, which it is not. That, to me, is the biggest myth that needs to be exploded.

**Soffer —**

I whole-heartedly agree with Kartik. This is a question I get a lot. People are so anxious about AI. Will it make decisions for us and take over? People get these ideas when they read the popular press. Then you go to the lab and see the gap between what they believe to be close to general intelligence and where we really are. That is what I think is overhyped.

Recently an idea that has received a lot of coverage in the popular press concerns a text generator—GPT-3 (generative pre-trained transformer), as it is called. The risks of this technology relate to what has come to be known as deepfakes. These things, on the surface, can seem very real—and that is why they are overhyped and get a lot of attention. It seems like a computer can write a poem or a news article. There have been a lot of articles about that.

When you look at what the computer is doing, you can see that it is not truly intelligent as we think of intelligence. But it is getting there, and it raises what I think are interesting questions about what intelligence is. Anybody who understands the technology behind it knows that it has ingested so much information that no matter what you ask, it can come up with what seems like a reasonable answer. That happens until you start asking it things it does not know. That is when you see that its response is nonsensical.

I do think something interesting is happening, and a lot of it is exciting. But on the other hand, it is far from what we would consider intelligence in the sense of understanding the rules or laws of the world

“  
**AI is good at narrow tasks such as reading X-rays and other such tasks in radiology. But to believe that AI is smarter than doctors makes it seem like general intelligence, which it is not**

”  
Kartik Hosanagar

and being able to reason or reach conclusions in a meaningful way. It is confusing, in that sense.

**Broudo —**

To turn now to the second half of the question, what are we not seeing? What do you think is around the corner, that people are not seeing yet and is getting lost?

**Hosanagar —**

For me, that would be the fact that when you look at AI, in practice, 90% of it is machine learning—and 90% of that is supervised machine learning. That is the idea that you have massive data sets which relate to what you are trying to predict. You learn from those data sets so that you can start making predictions. One of the exciting things that could happen is that there are other approaches as well. One of them, for example, is reinforcement learning. That is the idea that an AI system observes what happens and learns from it—and you do not need massive data sets to do that. The idea that you can create intelligence without massive training data sets is interesting. Reinforcement learning has been around for a long time—but in terms of industry applications and business settings in which it can be used, those are not as well understood. ☹



“

To me, the future developments that look promising are in what people are calling neuro-symbolic AI... That is what will help us go to the next level of being able to understand language as well as rationalize and reason. Adding knowledge and rules is where we want to go

”

Aya Soffer

Supervised machine learning is focused on its learning from the past, where lots of historical data exists. In contrast, reinforcement learning is about acting after observing what happens, and it allows us to learn from the past without training data. For example, imagine drug discovery. You could apply this concept to find drug molecules in situations where you do not have much training data. There are many interesting applications. The idea that you can learn even in situations where we do not have a lot of data from the past is super interesting.

**Soffer —**

To me, the future developments that look promising are in what people are calling neuro-symbolic AI. If we look at the history of AI, originally the concept was that we were somehow going to codify all the knowledge in the world. We had expert systems, we had logic reasoners, but that never really took off because it is impossible to codify all the world's knowledge and rules. Plus, it does not scale because it is hard to prove these things.

Instead, along came the paradigm of machine learning, which has turned out to be highly successful.

This was because, as Kartik has pointed out, it dealt with narrow tasks where, with enough examples, we could predict or classify things in a fairly good way. However, the issue we have today is that it is impossible to machine-learn everything.

First, this is because we do not have all the data; and second, because some things are much too intricate. This is also how we learn as humans. We learn some things through patterns: If you see enough images of cats, you know what a cat is. Second, we learn other things because we go to school, and somebody explains them to us. If you are taught that a cat goes “meow,” you know that a “meow” sound means a cat. Where we need to go now is to combine these two.

The neural architecture and machine learning can help us to learn better. We can use neural networks to perform logical reasoning on the data in the knowledge bases. That is what will help us go to the next level of being able to understand language as well as rationalize and reason. Adding knowledge and rules is where we want to go.


**Broudo —**

What challenges in AI and its deployment are unique to enterprises as opposed to AI in consumer applications? What critical differences should we be aware of when we think about AI in the B2B model versus the consumer model?

**Soffer —**

We at IBM work predominantly in the B2B model, so I will tell you what I have observed. The first thing to note is that B2B is B2B2C eventually. When you deal with businesses, you realize they have many issues around deploying AI. They care a lot, of course, about the outcome but they also care about the infrastructure. For example, questions can come up, such as, “How much it will cost to run AI?” or “Will those costs create enough value compared to the alternative of not running AI?”

A lot of economic considerations come up because running AI not only requires a lot of data but also computing power. It is important to understand the KPIs or key performance indicators around AI. It must deliver outcomes that are accurate, and it is equally important that AI must help a





company achieve its goals. The goals can range from functional ones—such as increasing sales—to non-functional ones, like the necessity to invest millions of dollars in computer farms. Those are some of the practical aspects.

Such considerations make AI deployment different than software engineering. How to deploy new software is relatively easy and now well understood; with AI, that is not yet well understood. We have this notion of AI life-cycle management. You train your models on data, then you test them—though no one knows yet how to test AI the way you test software. Unlike software, the performance changes if the data changes. In software engineering, the software does not change with the data. But in AI, if you have new data with other statistics, your models may no longer be relevant. These are some of the issues around deployment.

Another area I would highlight is explainability. For example, if a

borrower has been denied a loan [based on an AI recommendation], we must be able to explain why the loan was turned down. Companies are not happy with black-box AIs. We may develop more and more sophisticated black box neural networks, but people tell us, “No, please bring me back my rules and things that I can explain to people.” There is friction between what businesses want and feel comfortable with, versus the way AI really works. That makes it harder to deploy AI at the enterprise level.

**Hosanagar —**

Aya mentioned that B2B is often B2B2C. I would like to add a point from the end-user standpoint—whether that user is sitting in an enterprise or at home. An interesting question that often comes up is what it will take for that user to trust AI. Would a doctor be willing to trust the judgment of a diagnostic AI or would end-consumers be willing to trust a prediction from an AI system and apply it to their

“  
**How to deploy new software is relatively easy and now well understood; with AI, that is not yet well understood**  
”

**Aya Soffer**

life? That issue has not received as much attention as it deserves.

In fact, several studies in the social sciences and psychology show that people tend to have some algorithm aversion, especially when they see an algorithm fail. And no AI is a perfect system. It might on average be better than human performance, but an AI system can go wrong. How will people react when it fails?

Aya brought up explainability, but if you don’t even understand the system and you have seen it fail and you continue using it, what will it take for such an AI system to be adopted? Do we know if it is performing universally well for all people? Maybe it is beating lay people, but it is not beating experts. Some questions that relate to human psychology will start to matter a lot when it comes to designing the interface between humans and AI. It will also matter what information is shared with humans and what information is withheld to make it easy for them to understand and encourage adoption.

**Broudo —**

We turn now to natural language processing (NLP). Where is the field headed in the next couple of years? Why does it stand out in your mind as an area we should care about?

**Soffer —**

Let me begin with the second question about why we should care about NLP. Most of AI today is machine learning and analytics, but AI is about machines exhibiting

“  
**An interesting question that often comes up is what it will take for that user to trust AI. Would a doctor be willing to trust the judgment of a diagnostic AI or would end-consumers be willing to trust a prediction from an AI system and apply it to their life? How will people react when it fails?**  
”

**Kartik Hosanagar**

behavior that seems intelligent. If we look at human beings, what differentiates us from others and what makes us intelligent is the ability to develop human language to communicate. That is part of everything we do—whether we are talking with our family at home, in school getting an education, or at work. With a machine, that is not the case. I believe that as the field of AI matures, it will help machines to communicate with people using plain language. That will create the ability for every single person on earth to use computers and language. Humans can communicate through language by the time they are two years old, but computers still cannot do that.

The question that arises is, why has that been so hard? Where are we on the journey of trying to do better? Anyone who has used personal assistants [such as Alexa, Siri, etc.] or chatbots can understand that those programs are codified. They provide responses based on simple questions. That is where we need to go.

Technology in recent years—like GPT-3, which I mentioned before, which again stands for generative pre-trained transformer—has taken a big step forward. It can help with understanding words that are used in a certain context, but it does not help with the nuances of the language. That takes us back to the challenge of reasoning. That is what we will have to do to crack language. ➔



I do not think you can learn a language just by seeing more and more examples [of how it is used]. Reasoning in a meaningful way cannot be learned simply by crunching more and more data, but that will be necessary before computers can converse with us. Ultimately, natural language interfaces will do everything we do and completely change the way we interact with computers. For example, if you are a doctor, the computer should be able to recommend the next thing to do to help the patient. And it will also have explainability; the algorithm will be able to explain itself. It is all about understanding and generating language. Like AI, it will impact many applications in many industries.

**Hosanagar —**

I agree with everything Aya said, including what she said earlier about GPT-3 and the progress being made. It is an extremely exciting space—and one in which we will see a lot over the next few years.

**Broudo —**

In September *The Guardian* published an opinion piece titled, “A Robot Wrote This Entire Article. Are You Scared Yet, Human?” The newspaper claimed that the article—which was generated by GPT-3 — took less time to edit than many human op-eds. Is that why there has been such a buzz over GPT-3 compared with other instances of natural language generation? Where are the limits to what these technologies can do?

**Soffer —**

The reason people are paying more attention to GPT-3—compared to the previous version GPT-2, which was also amazing—is because of the amount of data it was trained on. Its full version has a capacity of more than 175 billion machine learning parameters. As a result, it can generate language that on the surface seems very natural. That is the reason why it has received so much attention. But if you were

to ask the algorithm to write about something new—such as COVID-19, before the pandemic was out there—it would not be able to write about something it did not know. It would not be able to write about a new vaccine, for example. It cannot do that. You would have to reprogram it with many, many articles about COVID-19 before it could write an article about COVID-19.

One strong feature of GPT-3 is that it has been able to pick up style. That is why people are getting excited about it. But several articles about GPT-3 have also pointed out its deficiencies. Often the words make sense, but the content does not always make sense. It is written in exceptionally good English, and it sounds like something that a highly educated person might write, but that doesn’t mean the content itself makes sense. It is interesting, and we will see where it goes. On the syntax, GPT-3 is excellent. On the semantics, it is still lacking. Still, it is a powerful NLP tool that will help us build better systems in the future.

**Hosanagar —**

In practice, creativity often means combining things in interesting ways. We should not assume that these systems are unable to come up with things that are novel and even creative. They will combine things in interesting ways. The scope of their creativity may be limited, but that is not in the same way that a lay person looking at it would rate it as being creative or not.

“**Ultimately, natural language interfaces will do everything we do and completely change the way we interact with computers**”

” **Aya Soffer**

A problem with these systems sometimes is that each sentence may make sense, but the paragraph may not make sense. At the same time, if you know what you are doing, you may be able to fine-tune GPT-3, give it very specific training data, and then tune it to produce the kind of result you want. If you think in a very narrow, targeted manner, you may be able to get it to do the kind of writing that seems almost human.

**Broudo —**

Kartik, since you are in the entrepreneurship space, how do you view investors’ approach to AI? Does every deal have to have an AI component? If so, does it drive a higher valuation? How well does the startup world understand AI, and is that different than what you have seen before?

**Hosanagar —**

AI is a big buzzword in the startup world. If a startup claims to have AI, that bumps up its valuation and increases the chances that the venture will be funded. As a result, lots of startups claim to have AI. When I said 90% of AI today is machine learning, I should have clarified and said that 90% of real AI is machine learning. The truth is that 90% of what passes for AI is not really AI—people claim that everything is AI if it touches data even slightly. A lot of that is going on.

That said, however, investors are starting to get savvy about AI. Some AI-specific venture funds have been created; some of them invest only in AI startups. As this happens,

I believe we will see more maturity in the market. Some investments are being made in pure-play AI startups—but those are becoming increasingly hard. Large players like Microsoft, Amazon and others are creating AI-enabled tools and giving them away almost for free as part of their cloud infrastructure. A startup that comes in with a horizontal AI application that can be applied across many industries will find it increasingly difficult [to generate revenues]. If Google or Amazon comes along with a free AI product as part of its cloud infrastructure, then it becomes extremely challenging for a startup in that space.

As a result, while there are some startups in the horizontal space, more are coming up in the vertical space and bringing AI to deal with a specific problem. For example, it could be AI for a personalized medicine application or for fraud detection in the credit card industry. We see a lot of that kind of activity. That is a little more defensible for the startup.

**Broudo —**

Aya, is IBM buying some of these startups? Do you see value in investing rather than

building? If so, what kind of companies are you looking for?

**Soffer —**

Generally speaking, many startups are building horizontal capabilities or tooling for the AI world. I agree with Kartik that they will find it hard to remain independent and grow that business. Companies want to run their AI on the cloud or the hybrid cloud, which is a combination of on-premises (or private) cloud and the public cloud. On the other hand, there are many ways a small company can innovate faster than a big company can. I do believe these smaller companies will be absorbed eventually by some of the larger ones. Eventually these capabilities may become part of the big platforms. The startups that will become large in AI will be those that focus on specific industry use cases.

**Broudo —**

Finally, as a last question: Investors can throw their money in places that change our world but one could say there is no clear governing authority. What should guide us in the absence of global standard rules?

**Soffer —**

Education is important; I would augment that with transparency. Regulation may help AI to become more transparent so people can make better decisions. In IBM, specifically, we are doing our best to pursue an idea—in partnership with other companies, of course—that we call AI FactSheets. If you think about nutrition labels for food, those were not there in the beginning. Over time, more and more regulations dictated that companies had to display labels on food packaging so that people could know what they were eating or drinking. Similarly, with AI factsheets, we will have a form that says you need to describe your model, how you trained it, how accurate it is, and things along those lines. That transparency, which is something that can be regulated, will let people know what is healthy and what is unhealthy in the ☹





## Governments must start getting savvy about how AI can be regulated without stifling what is innovative

” Kartik Hosanagar

consumption of AI—just as they do today with food consumption—so that they can make better and more informed decisions.

### Hosanagar —

We spoke about how AI can be transformative, and it is progressing at a rapid pace. We also discussed how we are in a world of artificial narrow intelligence and are inching towards general intelligence. We are going to have situations where AI can be used for good, but it can also be used irresponsibly.

Aya mentioned technologies such as GPT-3 and the fact that in the wrong hands, fake news articles can be produced at scale without any human beings being required to generate them. Photos and videos can be doctored to produce deep fakes. There is also the business of using AI for loan approvals or in the judicial system to make parole decisions. We do not need someone to have nefarious intentions for things to go wrong. All that is required is a slight oversight—and you may end up with a biased algorithm that makes discriminatory loan decisions that impact millions of people. It is not that humans are not biased; I do believe that, on average, AI will be less biased than humans. A biased judge might affect the lives of 200 or 300 people; a biased HR manager may make poor decisions about a few thousand people; but if an AI system is deployed to make decisions at scale, bias in those decisions may impact millions of people.

I do believe we need governance standards. The industry is participating in forums such as Partnership on AI to discuss best practices. My observation is that scientists from leading companies are coming together to discuss how to use AI responsibly. Still, when push comes to shove and decisions are made higher up in these organizations with a view to meeting quarterly targets, some of these conversations might not matter.

The governance frameworks should not be limited to companies self-regulating, in my opinion. Governments must start getting savvy about how AI can be regulated without stifling what is innovative. This will require participation by consumers, who will need to be educated about the technology and its risks. AI should be part of the curriculum in schools, so people understand what AI is and what it can and cannot do. For example, if they apply for a loan, they should know what assumptions have been built into an automated system. Or, if they read an article online, they should know how to assess whether the information they are consuming is truthful. Most importantly, education will need to change so that people know how to function in a world where AI is an active participant. ■



### About

**Dr. Aya Soffer** is VP of AI Technologies for the IBM Research AI organization focusing on natural language understanding and conversational systems and their application in customer care and other enterprise applications. In this role Dr. Soffer is responsible for setting the strategy and working with IBM scientists around the world to shape their ideas into new AI technology, and with IBM’s product groups and customers to drive research innovation into the market.

In her 20 years at IBM, Dr. Soffer has led several strategic initiatives that grew into successful IBM products and solutions in the Big Data and AI space including the original Watson system and more recently Project Debater. She has authored over 50 peer-reviewed papers and served as an invited speaker in numerous conferences.



**Professor Kartik Hosanagar** is the John C. Hower Professor of Technology and Digital Business and a Professor of Marketing at the Wharton School of the University of Pennsylvania. His research focuses on the digital economy, in particular the impact of analytics and algorithms on consumers and society, Internet media, Internet marketing and e-commerce. Professor Hosanagar serves as a department editor of Management Science and has previously served as a Senior Editor of *Information Systems Research* and *MIS Quarterly*.

Professor Hosanagar cofounded and developed the core IP for Yodle Inc, a venture-backed firm that was acquired by Web.com.