

Cognitive Technologies

AI and the revolution of work.



a purely AI-based call centre will likely not be very satisfactory and customers will continue trying to get the human operator on the line.

On the other hand, today we're also seeing more people working alongside AI. If you're not paranoid about what AI will do to work, then you're not really paying attention. Therefore, we have to constantly think about what the role of humans is, especially how they can add value to what AI can do.

The best results are achieved from humans collaborating with AI, rather than AI doing all the work. It will also be true for a while at least that if we want interesting and error-free content, we'll have to let humans take a pass at it. There have been some experiments, where people were given a choice whether to let AI-generated output be the final product or review the AI content themselves first. In one Massachusetts Institute of Technology (MIT) study on how its subjects used AI, 68 percent of them did not review or edit the AI-produced work at all. That's a bad sign, and we have to watch out. We have to encourage people to adopt a critical perspective on work produced with AI, so that they can figure out how to make it better.

For now, we have this environment where people and AI are going to be working with each other. I do believe that people who use AI in their jobs will generally be more productive and effective than those who don't. So, if you're a radiologist, you most likely won't lose your job to AI; you might lose it to another radiologist who uses AI. Not in the short run,

however, since there is a global shortage of radiologists!

It's therefore incumbent upon people who run organisations to make sure that their employees understand AI, use AI, and are critical in their application of this technology to their jobs, while ensuring humans are in charge of the final outcome. There are some companies that are already doing that. At PwC, for instance, AI has been introduced to all its employees.

What are some examples of successful use cases you have seen where GenAI has enhanced customer experience and service outcomes?

On the customer front, we've always believed what every business school around the world has advocated—you need to listen to your customers, understand them, and act on their inputs. In reality, customers are very diverse, and they send unwieldy messages that come through from different channels. As a result, reading and responding to them all is very labour-intensive. At the same time, we're inundated with content, and this is certainly even more so with us living in the attention economy, where getting people's attention for something that matters is becoming increasingly difficult. I think GenAI can really help in that regard.

In a consulting company that I co-founded, we worked with a retailer client on dealing with customer comments via email and social media, and so on, and realised that GenAI could do some things that typically could not have

Professor Tom Davenport, the President's Distinguished Professor

of Information Technology and Management at Babson College, speaks about how companies can integrate generative Artificial Intelligence (GenAI) into their operations while ensuring workforce adaptation and skills development.

How is AI transforming job roles and labour markets? How can companies integrate AI into their operations while ensuring workforce adaptation and skills development?

I try to be both empirical and optimistic about that issue, and so far, we haven't seen large-scale job losses yet, even though AI has actually been with us in various forms for 40 to 50 years now. In the long run, my guess is over the next five to 10 years, this situation is not going to change much, and we won't see massive layoffs. And in the areas where we do see layoffs, for instance, at call centres, the level of customer service that we'll get from

been done previously when dealing with multi-faceted comments. For instance, a customer commented that the meat from the deli was stale and shouldn't have been sold anymore; on the other hand, there was also a comment that the staff at the point of sale was very helpful and had returned the customer's money while complimenting the individual on his loyalty to the store. All that was in one single message. It turned out that GenAI, if given the right set of prompts about who is responsible for what, can determine that this positive comment should be forwarded to the store manager, but the negative comment is probably the responsibility of the meat department, so it should be

The best results are achieved from humans collaborating with AI, rather than AI doing all the work.

notified. GenAI would also suggest that perhaps there was a supply chain issue and would assign a probability rating to determine whether the supply chain manager should be informed.

Another company was using GenAI to examine conversations of staff who were calling customers to get them to settle their unpaid bills, and asking if they could work out a payment plan. GenAI was very good at that as well, and could even tell whether the call centre agent was following regulatory guidelines on how to treat customers. Eventually, the company figured out how to make GenAI assign a probability rating to assess whether the customer would ever pay, which



would help the company decide whether further efforts to call the person would be worthwhile. While these are all things that humans can do, we haven't really done them very well in most organisations, because they're so labour-intensive and require a fair amount of knowledge about how the organisation works. And it's not necessarily even an entry-level role. So, there are all sorts of possibilities out there for what can be done. Again, smart organisations will have humans review the messages first before they are communicated to customers.

What should C-suite executives consider when deciding how to use and measure the efficacy of GenAI?

This is a very big issue. It's particularly important for GenAI because most companies are really implementing it on the basis of productivity gains, and unless you measure it carefully, you're really not going to know this. In many cases, companies should do a controlled experiment or have a couple of different variations for the treatment group that does use GenAI in terms of the work processes they follow. Unfortunately, most organisations don't have the discipline to measure what they're doing in that regard.

There have been some efforts by academia thus far to measure the efficacy of GenAI. In some cases, the results show productivity gains while in others they don't. All organisations really need to look at measures like 'How many

customer messages have been dealt with hourly? What's the level of customer satisfaction?' If you're creating marketing messages, you would know what the outcome has been—whether people “click-through on it”. If it's something digital, it requires a fair amount of attention to measurement. We're not seeing a lot of that yet, and there's already beginning to be a small backlash to GenAI, with people noticing that it may not be yielding the productivity benefits that it should have.

What are some of the common challenges organisations face in adopting AI technologies and their strategies to overcome them? What practical advice would you have for businesses at various stages of AI maturity?

Some of the surveys I've done suggest that data is a big challenge, particularly in GenAI, where data is generally unstructured and typically in the form of documents, so you really have to carefully curate and manage it. Morgan Stanley, for instance, was working with OpenAI a couple of years before anybody was knowledgeable about ChatGPT. But even before that, many years earlier, they had realised that the quality of the documents on their intranet was not really what it should have been. They embarked upon a process of curating the documents and built an offshore capability of 20 people in the Philippines who would classify each document in terms of

how unique, accurate, and up-to-date it was, as well as how well it was tagged, and so on. So when GenAI came along, Morgan Stanley could pretty much quickly identify 100,000 or so documents to feed into a language model and effectively implement knowledge management with it, making important knowledge available to its financial advisors and their teams.

The survey that I did at the end of 2023 also suggested that about 80 percent of Chief Data and Analytics Officers agreed that they needed a new data strategy to deal with GenAI. And the majority had not done anything yet, so there is a lot of work to do in that regard.

Is companies' data ready for GenAI? I think in general, the answer is 'No'. There's also significant behavioural change that is going to be necessary to get people to use the technology in the right way. Furthermore, these people are generally knowledge workers, and if I had to take one lesson away from studying knowledge workers, it is that they don't like to be told what to do. Historically, they have had a lot of autonomy which they enjoy. That's kind of why they sought out the job. So being told that this is exactly the process you need to follow if you're going to use GenAI is probably not going to be appealing to many of them. Therefore there's a big behavioural change issue as well. And if you factor in that challenge and the data issue, that's a pretty considerable set of things that you have to get working well first in order to successfully deploy AI.



‘Citizen developers’, i.e., people with domain expertise but little formal computer science training, have become more common in recent times. How can organisations find the right balance between investing in state-of-the-art AI technology and reaping the low-hanging fruit presented by no-code AI development platforms?

This is my current area of research, so I’m quite attuned to it. I have a book that will soon be published on that set of issues. Already there were fewer and fewer barriers between non-technical humans and the ability to create systems of various types, and I think whatever barriers there are will also rapidly go away with GenAI.

This trajectory would result in start-ups that might truly embrace this low-code/no-code approach to the extent that the start-up profile might have very few technical personnel. I think we should train every potential entrepreneur to use those kinds of tools. The big area now in the low-code/no-code front for many organisations is the Microsoft Power Platform because they already have deals with Microsoft, and for a relatively small sum or incremental sums, you can get access to Power Apps like Power Automate or Power BI. And I think every student should probably be trained in those. This way, we will have fewer boot camps for entrepreneurs on ‘Here is how you need to develop all or most of the technological capabilities you need to

There’s already beginning to be a small backlash to GenAI, with people noticing that it may not be yielding the productivity benefits that it should.

get your start-up up and running’. Instead, there will be a big infrastructural development that will make GenAI even easier to use. But there will still be some barriers to succeeding as an entrepreneur—like whether your business model makes sense—although you can ask GenAI that question too! On the whole, technological capabilities will be much more readily available to entrepreneurs than they had been in the past.

What about the importance of a data-driven culture within organisations?

The good news is that GenAI seems to be creating more of a data-driven culture. Surveys I have worked on suggest that the percentage of organisations saying they have such a culture has doubled in the last year or two. But there are organisational leadership issues that often get in the way of developing this culture. Data problems have been there all along and we are creating data at a much faster rate than we can manage effectively. So there are these issues of what data management is about and how its role evolves. I’m very interested in these executive roles that manage technology—Chief Information Officers (CIOs), Chief Technological Officers (CTOs), Chief Digital Officers, Chief Data Officers, and so on. In some cases, these roles are going to have to merge.

Also, a survey I have just completed suggests that there are too many of such tech chiefs, and even the incumbents themselves are confused about the scope of their responsibilities relative to

other C-suite jobs. With GenAI, it’s a coveted responsibility, and you would think that Chief Data and Analytics Officers who have historically owned AI might be the owners of it. However, that is not the case in many companies. CIOs and CTOs tend to have more access to senior management, so they kind of grab these GenAI leadership roles in some companies. This is an illustration of the fact that we don’t have clarity for these roles, and there are too many of them. I call those in this new role of managing a lot of these functions the ‘super tech leaders’. They are the ones who are responsible for several of these areas and have specialists reporting to them, who deal with the details of that particular resource.

How important is it to establish robust ethical frameworks to guide AI development and deployment?

That needs to be an area of focus. While most organisations don’t really have much of a framework, I think the real need is not just for a framework, but for a process and, increasingly, even an automated or semi-automated process to evaluate AI development and deployment. Some really aggressive companies now have thousands of use cases. DBS Bank in Singapore is a great example where they have hundreds of use cases. As a result, having each one carefully examined by a human is rather time-consuming. Hence firms are looking for alternatives. For instance, Unilever has worked with a company in London called Holistic

AI to develop a semi-automated process to examine every proposed use case and comment on how the GenAI initiatives may be impacting transparency, bias, and the ways in which people might be negatively affected by them. And the process assigns a green, yellow, or red mark of approval, depending on whether the initiative is good to go, has potential issues, or should not continue, respectively. And it turns out that at Unilever, very few have come out red and need to be revised totally. Most of them need only a minor change.

My general feeling is humans are so biased and make so many decision errors that it’s relatively easy for AI to be better at it than humans. In some cases, that will be enough to go ahead with using AI; in some other cases, we may say it’s better than a human, but we still need to improve it. That is the way we feel about things like autonomous vehicles. Yes, they cause fewer accidents than humans do, but we want them to be almost perfect, which is a really high standard, and so far, nothing seems to be able to live up to that. [AM](#)



TOM DAVENPORT

is the President’s Distinguished Professor of Information Technology and Management at Babson College, the Bodily Bicentennial Professor of Analytics at the University of Virginia Darden School of Business, and a fellow of the MIT Initiative on the Digital Economy