



The potential impact of ChatGPT and the new AI on business

Paul Henninger – Head of Connected Technology

KPMG UK

kpmg.com/uk



You can't have missed the press coverage and social media chatter surrounding ChatGPT, the accessible writing tool from OpenAI.

Following a stream of developments from major artificial intelligence (AI) projects in 2022, ChatGPT proved a tipping point in the popular imagination. And despite some of the amusing answers it's produced, many people are wondering how tools like ChatGPT will be used by organisations.

Anyone can sign up for a ChatGPT account, and ask it a question, or to perform a task, with a simple, conversational prompt.

Here's an example:

User

Actually i'm wondering how to protect my home from burglary

ChatGPT

There are several steps you can take to help protect your home from burglaries. Some of these include:

1. Install strong, durable locks on all doors and windows.
2. Keep doors and windows locked, even when you are home.
3. Use a security system, such as alarms or surveillance cameras.
4. Keep your property well-lit, both inside and outside.
5. Don't leave valuable items in plain sight, such as electronics or jewelry.
6. Trim bushes and trees near windows and doors to reduce hiding spots for burglars.
7. Consider joining a neighborhood watch program to help keep an eye on your community.

Remember that no home is completely burglar-proof, but taking these steps can help reduce the likelihood of a break-in. If you have additional concerns, may be a good idea to talk to a security professional for more advice.

Ask a question on almost any topic, and ChatGPT will have a reasonable answer ready. You can ask it to write a song or give you a 5-part framework for a corporate digital strategy. On most general topics, like the one in our example, the output will be sensible. But on more specific questions, it might get a fair amount of detail wrong.

People have used ChatGPT and its related technologies to negotiate discounts on phone bills, dispense therapy to real-life patients, write Python code, poems, songs or novels, and to take (or cheat in) exams. Its AI, and large language models (LLMs) generally, produce amazing results.

As such, they could signal a shift in the way communications and businesses work.

But it would be all too easy to assume that it's time to make room for our AI overlords. A number of writers have, with some irony, written about how AI will put them out of business.

That sort of panic is a mistake. To understand ChatGPT's potential, let's look at how AI tools like it work, what they're capable of, and how businesses can use them.

Write an ad
for a luxury
perfume.

Write a piece
arguing for /
against Brexit.

Write a poem
about a starry
night.

**ChatGPT responds authoritatively
– in a sci-fi style that reminds you
of AI tech on Star Trek.**



What's under the bonnet?

The most recent generation of AI is based on LLMs. Interestingly, ChatGPT combines an LLM with an interaction layer that uses reinforcement learning – more on which later.

An LLM is a neural network model that uses unsupervised learning to predict outcomes. I won't explain how, because the fact is we don't actually know why it works – only that it does. Among the many AI models developed, LLMs are uniquely unexplainable.

Language models (as distinct from large language models) have existed for a while and can predict the next word or phrase in a sentence. They use different techniques than LLMs and have different applications – auto-correct is a common use.



Adding the 'large' element involves training the models on a large collection of publicly accessible electronic documents



Adding the 'large' element involves training the models on a large collection of publicly accessible electronic documents. That collection (or 'corpus' in AI terminology) comprises many petabytes of data – one petabyte being a million gigabytes. Training a model on such massive amounts of data allows it to learn about many topics, as well as language patterns.

So LLMs are 'large' partly because of the amount of data they're based on. But also due to the size of the models themselves. A few years ago, a complex model might have had a couple of hundred parameters; LLMs have billions. ChatGPT's underlying LLM has 175 billion parameters, and is training on something like 500 billion 'tokens' (words or word fragments).

The advances we've seen so far are largely a result of efforts to answer a single question: how can a model with that many parameters do something useful?

Large amounts of venture capital, human effort and computing power are required to train LLMs. OpenAI received \$1 billion from Microsoft; other LLM projects have come from companies like Google and Facebook. Sophisticated academic research has backed most of the development that's gone on. Given the range of possible applications, an entire industry is being built on making LLMs useful. And it's expanding rapidly: ChatGPT racked up a million users within a week of launch.

So why has this particular application become so popular, so quickly? It's partly because people from non-technical backgrounds can use it for a range of tasks. Many professionals, creators and writers have already tried it: they are the universe of users, customers and citizens that need to accept AI for it to have real impact.

But there's another factor. ChatGPT also has an analytics layer, comprising reinforcement learning built using feedback from humans (known as 'labellers' in AI-speak).

To create this, the 'labellers' gave OpenAI examples of what a "good" answer would look like. Then they ranked ChatGPT output for a particular prompt, from worst to best, with the results used to train a separate 'reward' model. Finally, this was used in a supervised exercise to create a policy which formed the logic that makes ChatGPT's user experience (UX) so good.



Making AI work

How do we know if an AI tool works? We could wonder whether it passes the Turing Test. This was created by computing pioneer Alan Turing to define the point at which AI attains a 'human' level of intelligence. An extremely simple version is whether AI can persuade a person that it's human. But a more practical benchmark would be whether interacting with AI is helpful and feels natural.

From that perspective, there are three critical success factors for the use of AI in a public or business context:

A good UX



AI results must meet a certain threshold to be useful.

For the better part of a decade, AI tools have outperformed clinicians in determining whether MRI scans show cancer. Others can predict whether an employee will be successful at a company from their CV for 20 years. But these applications failed to gain traction, because the users that would need to adopt them weren't convinced by the UX.



Perhaps an extension of the Turing Test should be whether AI tools feel too smart or bossy; and whether we can ask our own questions, rather than being told what to ask



01

Failing well



ChatGPT answers sound coherent and authoritative, even when some of the details are flawed. This is what we call a good "fail state".



Failing well can be more critical to AI adoption than succeeding (i.e. being accurate)



If users have a poor experience, even just once or twice, they'll quickly lose trust in the tool.

Good fail states vary between applications. Sounding plausible, but getting the finer details wrong, isn't a great fail state for an investment advice tool. In some cases, a good fail state may mean asking for more information, or allowing users to refine the output via a conversation. ChatGPT does this, as do some of the better image generation AI tools.

02

Ethical boundaries



The use of AI is fraught with ethical complications.

There are many examples of language models "learning" to be offensive, because they absorb offensive content into their corpus. And let's face it: social media is rife with bad patterns for them to learn from.

It's actually pretty difficult to get LLMs not to learn certain things – even though that's what their analytical and policy layers are for. This is a particular problem for image-generation models, which are capable of producing images in the style of specific artists. That's why some recently developed tools have introduced artist rights protections.

Consent is another key ethical consideration. A non-profit, mental-health platform ran into trouble for using AI to provide health counselling – without informing the patients that the content was AI-generated.

03

Work in progress



LLMs can be deployed to solve a wide range of problems – some of them pretty specific.



They can be modified to summarise and classify legal documents; respond to customer enquiries; assist expert advisors; and generate engineering and architectural drawings. Such applications require labelling to produce a good UX – but far less of it than previous generations of language modelling technology.

All the same, there's work to do before we can use the latest generation of AI for customer, employee, citizen and business interactions.

The starting point for that work is UX. While modifying models and working on analytics, we must think about users, interactions and processes. Users may range from employees and expert practitioners to customers, regulators and legal supervisors. How will each of these transition to using AI-powered tools?

Counterintuitively, limiting what models can do may have a more transformative impact, as users are likelier to reject more far-reaching change.

Start-ups will have the luxury of targeting users who are comfortable with change. And they can start small, making headway before regulation and enforcement catch up with the potential impact of their applications at scale. However, larger firms will face a UX challenge from the outset; that will need to be the AI pilots.



When developing new AI applications, there are four important questions to consider:

What do users want?

01

This doesn't mean making a tool do what users want it to. It means understanding how they see their role in AI transformation. That includes the people who will use it, and the end-customers who will experience it. For example, clinicians and patients for a medical application.



How valuable is the potential change?

02

This is a question for the leaders and investors sponsoring AI development. What are the expected costs and benefits of a new application? A few years ago, an AI tool was developed to process insurance claims. But as there were only three people involved in that process, it offered limited value.



What does good failure look like?

03

AI will sometimes get things 'wrong'. As we've seen, the crucial question is: what's the UX like in such scenarios? How can we anticipate and design good fail states for them?



What do we not want AI to do?

04

We must be explicit about the ethical, regulatory, legal and other risk boundaries we need AI to operate within, and use the technology to constrain it within them.

Finally, it's worth pausing to think about what to do if AI really does drive the change envisaged.

I asked ChatGPT to write a concluding paragraph for this paper, but the tool was "experiencing exceptionally high demand". So instead, it shared a limerick about its status:

**ChatGPT is surely the best
But its servers are put to the test
With so many users chatting It's no wonder they're lagging
But they'll fix it soon, no need to fret!**

I'd say that's a pretty good fail state. But it's also a warning. Developers that offer the right outcomes and experiences will need to invest in some serious server capacity.

Contact



Paul Henninger

Partner, Head of UK Connected
Technology & Global Lighthouse
KPMG in the UK

LinkedIn: <https://www.linkedin.com/in/pahenninger/>



: Paul.Henninger@kpmg.co.uk



: +44 (0)7557 578902

kpmg.com/uk



The information contained herein is of a general nature and is not intended to address the circumstances of any particular individual or entity. Although we endeavour to provide accurate and timely information, there can be no guarantee that such information is accurate as of the date it is received or that it will continue to be accurate in the future. No one should act on such information without appropriate professional advice after a thorough examination of the particular situation.

KPMG Law is part of KPMG LLP, a multi-disciplinary practice authorised and regulated by the Solicitors Regulation Authority. SRA ID: 615423. For full details of our professional regulation please refer to 'Regulatory information' under 'About' at www.kpmg.com/uk

© 2023 KPMG LLP, a UK limited liability partnership and a member firm of the KPMG global organisation of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved.

The KPMG name and logo are trademarks used under license by the independent member firms of the KPMG global organisation.

CREATE: CRT146566A