



# The Four Essential AI Use Cases



While ChatGPT hype is relatively new, conversations around the power of neural networks and machine learning certainly aren't: back in 2016, The New York Times published a piece on [The Great AI Awakening](#). It described how although neural networks had existed for decades, they had never quite found widespread commercial use – until now. In 2017, Google published the now famous [Attention is All You Need](#) paper, which introduced the revolutionary transformer architecture.

But despite those publications, our 2017-selves would still be shocked by all the progress we've made in AI by 2023. And if the next six years are anything like the past six have been, we are going to be living in a very different world by 2030.

## So what's changed?

Well, there have been many advancements, including improvements in GPU hardware, the introduction of powerful pre-trained transformer models like BERT, a cambrian explosion of the AI open source community on [HuggingFace](#), and new breakthroughs in biology like [AlphaFold](#). But the most important advancement of all – and the reason that you're reading daily about AI in the press has been the rise of **generative AI**, and especially **large language models** like GPT-3 and ChatGPT.

"Generative AI" refers simply to AI models that generate things, like text or images. It's not new, per se: generative adversarial networks – a predecessor to image generation models like Midjourney and Stable Diffusion – were invented by Ian Goodfellow in 2014. And autoregressive language models like GPT-1 and GPT-2 have been around for even longer than that.

What's new is the scale of these models, as well as their training procedure: **increasing the size of these models has made them far more intelligent and flexible, but perhaps most importantly it's largely obviated the need for large training datasets.**

For example, if you needed a model to determine a user's intent when speaking to a chatbot in 2017, you would have had to train a model on tens of thousands of example messages, which means that you would have to label tens of thousands of messages – which is time-consuming and expensive. But in 2023, you can simply "prompt" GPT-3 on the task at hand, describing it as you would to another human, and GPT-3 will simply perform the task based on your instructions.

**The ability to simply prompt rather than train, along with huge upgrades in intelligence and coherence, makes LLMs a game changer for businesses and consumers.** It makes possible things that would have been considered science fiction only a few years ago. Used properly, it has the potential to create enormous value for the economy and, more importantly, genuinely improve peoples' lives. But for that to happen, we need to understand AI: its strengths and limitations; how to use it responsibly and effectively.

In this guide, we'll dive into what business leaders need to know about essential AI use cases and how AI will affect their organization.

## Four Essential Use Cases

Talking about use cases for AI in 2023 is a bit like talking about use cases for the computer in 1970. If, in 1970, you had tried to explain the potential use cases for computers, you might have come up with examples like performing mathematical calculations on large data sets or setting up a database for a company's inventory.

But, of course, that would have sold computers far short: we use them for much more than just doing math or keeping track of inventory. Computers suffuse every aspect of our lives, and today we use them for many things that in 1970 would have been completely unimaginable. AI seems likely to represent a platform shift as significant as personal computing, mobile computing, or the internet, in that it will come to suffuse every aspect of our businesses and lives just as those technologies did.

So, while identifying concrete use cases feels almost silly or limiting, they are an important place to start — especially as many organizations are struggling to understand where this technology fits in and what it can be used for. Further, separating out different use cases allows us to make better decisions about risk, because depending on the use case the risks are very different.

**For each use case, we'll outline example applications, the specific technologies, and algorithms it relies on, how human involvement works, and what's different now that we have generative AI.**

01

# AI as a Creative Assistant



**The first use case is using AI as a creative assistant. In this case, AI is helping a high-level creative employee by providing suggestions, retrieving information, and supporting output. Some of us already do this frequently in the form of [GitHub Copilot](#) or ChatGPT.**

For example, for a product manager, Copilot can help autocomplete code or write code based on prompted instructions. ChatGPT can serve as a sort of replacement for StackOverflow, helping debug things and think through novel problems (much more quickly than by reading documentation). These two products are already leading to increased productivity in these fields – and more.

Copilot is a special creative assistant for software engineers, but in the future, every industry and application will likely have some form of creative assistant embedded into it. There are already a number of other AI assistant products emerging for other industries such as [Jasper](#) and [Copy AI](#) for copywriting, or [UIZard](#) and [Galileo](#) for UI design. Additionally, many software products like [Notion](#) and Microsoft Office are beginning to embed AI assistants directly into their products.

This year, we will likely see the emergence of other industry-specific creative assistants for fields like law, medicine, entertainment, advertising, academia, and finance. Here are some specific examples:

- Legal creative assistants may help lawyers write briefs and contracts by quickly surfacing and summarizing relevant cases. Often the assistant will write the first draft and the lawyers will tweak as needed.
- Medical creative assistants could help doctors to diagnose patients based on a written set of symptoms or to write correspondence to insurance companies like [this person](#) already is doing.
- In entertainment, creative assistants might help screenwriters generate new ideas for TV episodes and even write the first drafts of the dialogue.
- Advertising creative assistants will help write first drafts of commercials or other ad copy based on a prompt. They'll also be able to create visuals or even entire videos.
- Academic creative assistants may help professors write papers by surfacing and summarizing relevant academic literature.
- Financial creative assistants could help bankers fill out complicated Excel formulas in their models and summarize relevant market data.

These applications are really only possible because of generative AI and LLMs. With the exception of [Clippy](#), this category of product essentially didn't exist before 2022, but since then we've seen an explosion of new startups. Before LLMs, AI models simply were not smart or flexible enough to provide much help in complex creative domains.

When using AI as a creative assistant, the human is still very much in control. The AI will make suggestions, surface/summarize relevant information, or try writing a first draft. But the human — usually a highly capable knowledge worker — ultimately has the final say. The human gets to accept, reject, or tweak the AI's suggestions.

**Overall, this use case has the potential to make knowledge workers much more efficient. It augments workers, rather than replacing them, with minimal risk.**

02

# AI as Business Process Outsourcing



**The second use case involves using AI to automate business processes. This is similar to how many firms automate certain rote, highly repetitive processes to other firms and other countries (hence “business process outsourcing”).**

A large number of business processes will be candidates for automation in the coming years. Many of them are esoteric tasks that are specific to each individual company, making it difficult to predict exactly what this will look like.

However, a few specific examples might include...

- Using AI to identify and fill performance gaps in [customer support](#)
- Medical coding, which is the process of translating diagnoses and procedures into standardized medical codes and costs the US economy more than \$400 billion per year, according to a [2016 study](#)
- AML/KYC (Anti-Money Laundering / Know-Your-Customer), which is the process of financial institutions doing due diligence on new and existing customers

Of course, it's been possible to automate tasks using AI and machine learning for many years now, and often older ML models are the right tool for the task. However, LLMs significantly raise the bar for what can be automated while simultaneously lowering technical barriers. LLMs are not just capable of writing: they can also reason, predict, and classify. This means they can be used to automate a wide variety of tasks, not just ones that involve writing prose.

What's challenging about outsourcing a process completely to AI is that there is no longer a human directly in the loop. This means the requirements for accuracy become much higher, and additional safeguards are required to monitor and retrain the model over time. One of the best ways to mitigate risk is to **fine-tune** or **supervise** the model, which means training the model on many thousands or millions of examples of the task, as opposed to using the zero-shot learning, where you merely give the AI written instructions. The more examples you have, the more accurate the model becomes, minimizing the risk of mispredictions.



But the risk is never zero. In evaluating this risk, however, it's important to always ask "What's the alternative?", which is usually to have humans continue doing the task. Humans have the disadvantage of being both costly and slow, but if they are a more accurate or safer option then that cost is justified. However, humans are not always a more accurate or safer option.

Therefore, the question should not be "Can AI do this perfectly?" but instead "Does AI improve the status quo?" and in many cases the answer will be "yes".

The lowest hanging fruit here will be highly repetitive back-office work in the information economy. Importantly, AI still struggles with applications in the physical world, as evidenced by the surprising lack of adoption of self-driving cars in the last ten years. It turns out that getting AI to interact with the physical world in a safe, predictable way is extremely challenging.

It will be a very long time before we have AI plumbers, chefs, or farmers. But AI medical coders? Or KYC analysts? Or paralegals? These are right around the corner, and it will behoove every business to automate this type of rote, manual work. Not only will it save enormous amounts of money, it will free your employees from some of the most thankless, unfulfilling tasks. This form of AI is likely to replace many jobs — we don't do anyone any favors by pretending it won't, but they are jobs that, ultimately, we will not miss.

03

# AI as Digital Experience



**A third use case is using AI to create or augment digital experiences. (We're using "digital experience" as a broad term to refer to any touchpoint between a business and its end users.) Essentially, AI is a medium through which the user interacts with a business or uses its services.**

This use case is already well-established – many people have been interacting with AI-enabled digital experiences for years now.

We use an AI-powered search engine every time we interact with Google. We use AI-powered recommendation engines to browse movies, songs, and products on Netflix, Spotify, and Amazon. We use AI-powered conversational agents like Siri and Alexa to look up information and perform tasks. **AI is increasingly the medium through which we interact with digital businesses.**

When it comes to AI-powered digital experiences, we usually aren't trying to replace a human but rather provide an experience that fundamentally can only be delivered by software and AI. The human's role – to the extent that there is one – is more to monitor the AI's behavior and curate the content that it has access to.

AI-powered digital experiences have existed for nearly a decade now, but generative AI has the potential to enhance them significantly. Historically, building AI-enabled digital experiences was only possible for the largest, most advanced tech companies, advances in AI are making it possible – even necessary – for more companies to deliver these experiences.

This takes many forms, such as...

- A high-tech company providing a [virtual assistant](#) that can help users navigate their documentation and answer complex questions about it
- A retailer providing personalized recommendations for complementary products, as well as fast semantic [search](#) to help users find the precise product they're looking for
- A health system offering a virtual assistant that can help customers identify the right doctor to treat their symptoms and get an appointment scheduled

As AI gets more powerful and widespread, consumers' expectations will rise, and they will no longer tolerate clunky, outdated websites and chatbots. They will expect to be served relevant, dynamic content and instant answers to their questions. If they don't, they'll take their business elsewhere.

04

# AI as Prediction and Forecasting



**The final use case is one that you won't read about as often but is extremely valuable: AI can be used to predict the future. Astute readers will point out that all AI is predictive: LLMs are *predicting* the next word in a sequence, diffusion models are *predicting* the ideal set of pixels for an image, and recommendation engines are *predicting* whether a user will like a product.**

In this case I'm referring specifically to predicting business outcomes in the real world. For example:

- Predicting consumer demand for a specific product based on historical trends
- Predicting what the stock market will do tomorrow based on a variety of factors
- Predicting how likely a person is to repay a loan based on their payment history
- Predicting the weather tomorrow based on IOT data

You might think of these things more so as “statistics” than “machine learning” or “artificial intelligence,” but, as it turns out, those are all the same thing. **AI and machine learning *are* statistics, applied at an extraordinary scale and to problems that we don't traditionally think of as involving numbers.**

This use case is the oldest of the four. We have been using math and statistics to predict the future for many decades now. What's different today is that the algorithms can handle much more data than they previously could.

Historically, we would feed predictive models only a carefully curated set of signals or “features” that were known to have strong statistical relevance. (For example, if you wanted to predict the price of a house, you'd want to use square footage as a feature.) We might use ten or twenty of these data points to train a simple predictive model.

But, over the past ten years, deep learning and large neural networks, combined with other innovations in data processing, allow us to train much larger models with many more features on massive datasets. These models are able to learn complex relationships and representations that would have been impossible to learn with more primitive statistical methods.

Here once again the whole point of the model is to outperform a human, and usually these models do so handily. In fact, these models excel the most in the exact situations where human intuition fails. If you asked the average person to predict movements in commodity markets based on historical time series data, most surely couldn't. This is a job for an algorithm. An algorithm can process far more data and observe statistical patterns that no human could.

This use case is probably the most unaffected by generative AI. In general, generative AI does not move the needle very much on our ability to produce statistical forecasts. For the most part, older ML methods like recurrent neural networks and even transformers work better here. But there are some exceptions: for example, Deepmind recently released [Graphcast](#), an algorithm that predicts the weather with state-of-the-art accuracy using generative AI. Perhaps generative AI will revolutionize this field as well.

## Conclusion

The use cases we've discussed are by no means exhaustive, and many applications don't fit neatly into a single category. (For example, facial recognition software on your iPhone.) But this framework helps us to frame AI applications for your business, for example:

AI can be used to assist creative knowledge workers such as programmers, lawyers, doctors, teachers, or screenwriters. You should adopt this sooner rather than later because there's very little risk and AI can be extremely helpful.

AI can be used to automate entire business processes, wholesale, but this is riskier and requires high accuracy and additional monitoring/safeguards. But sometimes the risk is worth taking, and you should compare AI to the status quo, not to perfection.

AI can be used to create immersive, helpful digital experiences. The human's job here is to curate content and tweak the algorithm's behavior. This is relatively low risk, and consumers will increasingly demand AI-powered experiences of all companies — not just Google and Netflix.

AI can be used to forecast the future based on historical trends and other data. This is a valuable application, but relatively little has changed in this arena.

By using AI to collect and organize knowledge across the enterprise you can deliver scalable, seamless digital experiences.

**Ready to learn more? [Request a demo today.](#)**



Yext (NYSE: YEXT) helps organizations build digital experiences across any channel on our open and composable platform. The Yext Digital Experience Platform collects and organizes content to deliver AI-led experiences for any organization's customers, employees, or partners. For more than 15 years, thousands of companies worldwide have trusted Yext to create seamless customer experiences at scale across search engines, websites, mobile apps, and hundreds of other digital touchpoints.

Learn more at [yext.com](https://yext.com)