


THE APPLICATION OF LARGE LANGUAGE MODELS IN SERVICE MANAGEMENT



[ChatGPT](#), OpenAI's brainchild that launched in November 2022, seems to have generated more interest in the last couple of months than any other technology or trend has in the last couple of years. The industry titans have certainly taken notice, with Microsoft putting its weight behind ChatGPT both [financially](#) and [technologically](#) while [Google announced Bard](#), its own artificial intelligence (AI)-powered chatbot.

Based on GPT-3 and powered by humongous data sets backed by heavy computing requirements, some have claimed this as one of the biggest technology-based disruptions that will fundamentally alter the economics of information, while others have pointed out the limitations of GPT-3-like large language models (LLMs), and then there are all sorts of opinions in between.

Where did all these enthusiasts come from suddenly? LLMs are not the new kid on the technology block.

One LLM differentiator is that it can be used to consider very large volumes of data such as text, and understand the overall context, versus natural language processing (NLP) that only looks at the immediate context between words.

BMC has been experimenting with LLMs for years now. It might seem like a new entrant into the market, but there's still a way to go before it reaches maturity. Will it serve as a lamp post that can be used not only for support but also for illumination, or will it fade out alongside other hyped,-but-failed technologies? Only time will tell.

In this post, we explore three leading use cases from a service management standpoint that makes

us believe that the democratization and widespread use of AI has just begun. The traditional tradeoff between “richness” and “reach” might be a thing of the past. Generative-AI has the power to deliver both—to all participants and on demand.

1. **Ticket text summarization**

In the service management world, we deal with large volumes of text that could be part of activity or work logs, chat conversations between agents and subject matter experts (SMEs), meeting transcripts or notes, feedback forms, knowledge articles, or even resolutions submitted for hundreds of thousands of tickets. When we use an LLM that is trained on more than one billion parameters, it can lead to a rich summary of information collected from piles of data that can now be presented to the service desk agent.

The insights generated from that summary would then help them to arrive at the root cause or the solution within minutes, thus saving valuable time that could now be spent performing value-added tasks. Also, summarizing text from user feedback forms can succinctly point service managers to areas of improvements and can be used by them for training purposes. This would ultimately lead to increased efficiencies in the service desk, a streamlining of the operating model, and eventually, a positive improvement in end-user experience.

1. **Generating knowledge articles**

Knowledge management is important to the service desk because it allows teams to capture, retain, and share knowledge that—when used effectively—can decodify the tribal knowledge that lies with a few specialists. That makes it easier to upskill and enable the workforce to make the right decisions throughout the service lifecycle. ITIL[®] 4 also defines knowledge management as one of the central processes responsible for providing knowledge to all other IT service management functions.

So, though this is of high importance, it leads us to ask a question: How many times does a technician document the solution(s) that they provide when resolving a ticket?

The response would vary from organization to organization, but what we have seen is that not all solutions become a knowledge article. In fact, most of the time, only a small fraction of them gets converted, which creates a knowledge gap between an average agent and the subject matter expert. This is where generative-AI technology can assist with the heavy-lifting.

By looking at the combination of incident and problem data such as activity logs, solutions provided, resolutions entered, and identified root cause, the model can generate drafts of knowledge articles that can be submitted for review by the agents or experts. Instead of starting from scratch, agents and experts can now review and revise the contents of the generated knowledge article within minutes and submit the article for approval and publishing. Undoubtedly, automatically generating knowledge articles using AI technology will lead to better and greater dissemination of information across the organization.

1. **Data cleansing**

Data is the holy grail of machine learning. But predictions based on bad data usually results in a case of “garbage in, garbage out”. If the quality of outcomes must be improved, then the underlying data needs to be parsed, standardized, matched, and cleansed. For example, spelling mistakes made while logging an incident can result in matching wrong knowledge articles or no knowledge articles

being recommended.

Properly trained LLM models can provide the right level of inputs for data cleansing. Some examples could be during a product upgrade where all form data is spell-checked; pre- or user-populated content is verified; or jargon and abbreviations are converted into layman language. This can ultimately help secure the right data and context to ensure that a machine learning (ML) project is using cleansed, standardized, and trusted data.

The “art of the possible” with LLMs also extends to use cases where the search experience can be improved on a vast underlying dataset; with chatbots that can combine data from within and across the organization; or even by developers who can convert screens into documentation. In conclusion, we would like to say that a new era of AI has dawned upon us. This one is super-exciting and the one that we will be watching with bated breath.