


TAKING CONTROL OF AI



**CUSTOMIZING
YOUR OWN
KNOWLEDGE BOTS**

AUTHORS

Nicholas Johnson, Michael Papadopoulos, Foivos Christoulakis,
Gregory Brown, Anastasiya Tikhonova, Eystein Thanisch

For decades, organizations have struggled to connect employees with the internal knowledge and insights they need to make better, more informed, and timelier decisions. Finding the right information, at the right time, among the increasing volumes of structured and unstructured corporate data, and delivering it in the right format is business-critical. However, it is often like looking for a needle in a haystack.

AI can finally overcome this challenge. Essentially, the ability of Large Language Models (LLMs) such as ChatGPT to create bespoke content is game-changing when combined with Natural Language Processing's (NLP's) capacity to find, process, and format information. However, using completely open LLMs as the engine for a critical business process such as enterprise knowledge management may give executives pause for thought.

Based on first-hand research and practical pilots, this article sets out how companies can harness the huge benefits of AI to manage knowledge across the enterprise, without having to take the risks inherent in relying entirely on an open-access LLM.

THE DRAWBACKS OF LARGE LANGUAGE MODELS

The explosion in the use of ChatGPT and other LLMs promises to transform business and society. Better sharing of knowledge and information is a key use case for business, with clear benefits for companies and employees. However, simply introducing a third-party tool such as ChatGPT into an organization to underpin intelligent knowledge sharing through AI bots has multiple drawbacks:

- It exposes your corporate data to the provider of your LLM (OpenAI in the case of ChatGPT). While this is acceptable for generic data, it risks leaking proprietary knowledge to the world. In a recent example, Samsung banned staff from using LLMs after sensitive data was uploaded to the ChatGPT platform.

- Due to their size, bigger LLMs have a training cut-off period, and don't "know" about anything that has happened since then. ChatGPT, for example, is only aware of data up to September 2021.
- LLMs have been shown to suffer from "hallucinations," confidently providing answers that are fundamentally inaccurate. And, given that they are trained on the internet, the information they provide may be based on untrue data, a risk that could increase over time as AI bots generate increasing amounts of that data themselves.
- In a fast-emerging market, regulators are struggling to catch up with technology. This introduces the risk that if you embed a third-party LLM into your organization, regulators may subsequently decide to impose conditions on its use. For example, the Italian government temporarily banned ChatGPT in April 2023, citing privacy concerns. Many countries are moving towards regulating AI in some format.

TAKING CONTROL OF AI

Rather than being reliant on third-party tools, organizations can instead implement their own LLM. This should be built on one of the numerous commercially exploitable base models and focused and trained on the organization's own internal data. This provides a significant advantage in terms of privacy and security since data is retained within the organization's infrastructure.




FIGURE 1: HIGH-LEVEL BOT ARCHITECTURE

It also avoids issues that could arise from the sudden imposition of regulatory restrictions. For instance, since the organization is fully aware of the data being used to fine-tune its model, it can satisfy any “right to forget” obligations by deleting the model fully and then retraining it without the infringing dataset.

Smaller base models also overcome training cut-off periods. They can be continually trained and fine-tuned by adding data from the organization itself using a variety of techniques such as Low Rank Adaptation (LoRA) or Quantized LoRA (QLoRA). This provides sufficient depth, as information is focused on the organization itself – rather than extraneous general knowledge that is irrelevant to normal business operations.

These factors all mean local internal models are best suited to corporate needs. They can be used to create powerful knowledge bots, intelligent assistants that help employees in specific areas. Essentially, these bots should be T-shaped, with deep knowledge and expertise in a single discipline. This means you may need multiple bots, each specializing in a particular subject.

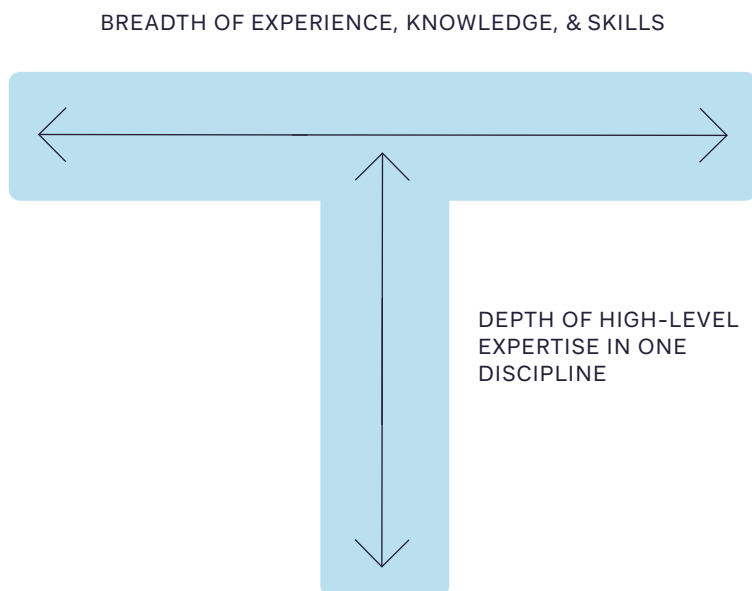


FIGURE 2: T-SHAPED SKILLED AI MODELS

These locally running models don’t require wide knowledge bases and can focus on specific tasks, unlike larger models that need extensive knowledge. This not only improves their performance, but also makes them easier to train, maintain, and update, effectively providing smaller models with the ability to carry out the specific task that you’ve trained them for.

By being trained on specific datasets that incorporate the organization’s proprietary knowledge and industry-specific information, knowledge

bots will acquire a deeper and faster understanding of relevant terminology, jargon, and concepts. This enhances their search capabilities and delivers more accurate and relevant results. It also makes them much less likely to generate misleading or incorrect information, a major advantage in situations in which accuracy is paramount, such as medical diagnosis and legal advice.

A PERSONALIZED APPROACH

Custom-trained knowledge bots can also deliver customized responses that cater to the unique requirements of each user in a way that generic LLMs can't. This personalization enhances user satisfaction and encourages greater engagement, fostering a strong

CUSTOM-TRAINED KNOWLEDGE BOTS CAN ALSO DELIVER CUSTOMIZED RESPONSES THAT CATER TO THE UNIQUE REQUIREMENTS OF EACH USER IN A WAY THAT GENERIC LLMS CAN'T.

“personal” partnership between employees and knowledge management systems. Users can interact with a bot using natural language queries, creating a conversational interaction that improves the experience and eliminates the need for complex search queries, technical commands, or

detailed training. If the user doesn't understand the initial answer, they can ask for it to be explained in a way that better meets their requirements.

As bots build their knowledge of users' preferences and previous interactions, along with the context within which they operate, they can act as intelligent information filters to ensure time and attention are directed onto not just the most productive tasks, but also the most important ones. For instance, a knowledge bot integrated into a project management system could identify potential roadblocks well in advance and provide proactive recommendations to keep projects on track.

DEMONSTRATING THE PRODUCTIVITY GAINS FROM AI BOTS

One of the primary benefits of an effective knowledge bot should be that it improves productivity by providing the right information more rapidly than conventional approaches. However, given the rapid rise of generative AI, little published material exists to quantify the benefits when used internally within organizations.

Here at ADL, we therefore undertook a six-week pilot to determine the potential productivity improvements of AI used internally within an organization. We created a knowledge bot focused on the Atlantis Terraform automation tool, an expert solution that requires very

specific knowledge to operate. The bot was trained with all publicly available data (such as vendor product manuals), along with internal troubleshooting knowledge from users.

The bot was then released for pilot use by a team of 18 Service Reliability Engineers, all of whom use Atlantis significantly in their roles. The following methodology was used:

- Team members marked the bot’s answers as Complete, Incomplete, or Wrong.
- They then marked the answer as Useful or Useless.
- In the case of Useful answers, engineers estimated the time saved in five-minute increments.
- For Useless answers, 15 minutes were subtracted from the time-saving total. This was an estimate of the time wasted accessing the bot, forming the question, and waiting for the reply.

On average, the bot was asked 300 questions per week. Approximately 45 percent of responses were marked as Complete, whereas the rest were marked as Incomplete. Only a single answer was marked as Wrong across the entire experiment.

In an engineering setting, even Incomplete answers can be Useful – as they can help direct people to a solution to their problem. This led to




FIGURE 3: RESULTS OF SIX WEEKS OF BOT Q&A

70 percent of answers being marked as Useful, indicating an overall time saving.

Aggregating the results showed that, on average, approximately 40 hours of effort were saved each week – a nearly 10 percent productivity gain or the equivalent of employing one extra team member. In addition to delivering an immediate return on investment, this productivity gain is likely to increase over time, as:

- The system was new, and engineers were not required to use it.
- Processes were not changed to incorporate the bot.
- Users received minimal training – they were just asked to type their questions into the bot.
- Employees were still getting used to the bot and finding their own ways to make it work for them.

Additionally, the pilot involved a small, 18-person team of experts. Productivity gains will become greater as the use of knowledge bots is scaled up. For instance, it would add the equivalent of 10 people on a team of 180, leading to the potential for reallocating resources to do more with a lower headcount.

UNDERSTANDING THE BENEFITS BEYOND PRODUCTIVITY

Besides the productivity benefits they bring, knowledge bots can deliver advantages in other areas:

- They **promote collaboration** across organizations by offering employees from different departments fresh insights through a company-wide knowledge base. This drives interdisciplinary collaboration, leading to development of new ideas and solutions that might not have emerged within isolated teams.
- Multilingual knowledge bots can **overcome language barriers** and promote diversity and inclusion within organizations by allowing employees from different countries and cultures to exchange knowledge on equal footing.
- Bots can serve as **virtual coaches or mentors**, providing personalized guidance and support to employees on their professional development journey.

THE CHALLENGES TO ONGOING SUCCESS

For all the benefits custom-trained knowledge bots bring to organizations, they also introduce challenges that need to be understood and overcome, especially around the ethical implications of relying on AI to generate and disseminate knowledge. For instance, as with all AI, organizations need to be aware of the origin and accuracy of training data, even when generated in-house, and the consequent potential for misinformation and bias.

Transparency must therefore be a priority, especially in sensitive domains in which bot-assisted decisions must be documented and justified. This means humans need to be in the loop to provide accountability and ensure organizations operate within ethical boundaries and adhere to legal regulations.

The question of intellectual property rights and ownership of generated content must also be considered. As bots aggregate and synthesize information, sources and respect for intellectual property must be attributed, based on clear, established guidelines and procedures.

Additionally, employees may initially resist using bots for fear of losing status or even their jobs. Introducing custom-trained bots

will therefore require a well-thought-out change management strategy. Involving employees early on, clearly communicating the benefits of using bots, and providing training on how to get the best from them can do much to reduce concerns, as can encouraging employees

to continuously explore and experiment with their new personal assistant.

Finally, corporate knowledge changes rapidly. If custom-trained knowledge bots are to maintain their long-term usefulness, they need to be regularly fed with new data and insights so they evolve with an organization's needs. Techniques such as transfer learning, domain adaptation, and active learning can be employed to improve the bots' ability to learn and adapt to new information. Organizations should also put in place continuous monitoring and feedback loops to help identify and rectify any inconsistencies or inaccuracies in their outputs. Quality assurance processes, including human oversight and validation, ensure the generation of reliable and accurate information.

THE QUESTION OF INTELLECTUAL PROPERTY RIGHTS AND OWNERSHIP OF GENERATED CONTENT MUST ALSO BE CONSIDERED.

INSIGHTS FOR THE EXECUTIVE – ACHIEVING COMPETITIVE ADVANTAGE THROUGH CUSTOMIZED AI BOTS

AI-powered bots are poised to transform knowledge management within organizations, delivering greater efficiency, productivity, and customer satisfaction, along with hyper-personalized learning experiences for employees.

Investing in custom-trained knowledge bots therefore has the potential to create a significant competitive advantage for a company by streamlining processes, improving operational efficiency, automating time-consuming activities, and accelerating workflows, as shown by our initial research. As technology continues

THANKS TO AI, WE WILL ALSO SEE GREATER COLLABORATION BETWEEN HUMAN EXPERTS AND KNOWLEDGE BOTS, RESULTING IN THE CO-CREATION OF KNOWLEDGE.

to evolve, these benefits will undoubtedly expand, transforming the way organizations access and use knowledge.

These productivity improvements can help create a significant competitive advantage in a dynamic business environment, as does

the faster, more accurate decision-making capability that knowledge bots underpin. However, in the future, the greatest differentiation may be from the bots' ability to generate insights and uncover valuable patterns from diverse and high-quality data.

Thanks to AI, we will also see greater collaboration between human experts and knowledge bots, resulting in the co-creation of knowledge. This combines the bots' capacity for data analysis with the unique insights, creativity, and contextual understanding of human experts.

As knowledge bots evolve, we could see bot-to-bot interactions, with bots collaborating and exchanging information as they solve complex problems. This would lead to networked intelligence within and between organizations.

To gain the benefits of AI within corporate knowledge management, CEOs should focus on:

- Understanding the AI market and going beyond tools such as ChatGPT when setting strategy and embracing opportunities. Look to bring AI development in-house using available base models
- Investing to develop the right skills and capabilities to build and optimize knowledge bots
- Focusing on foundational enablers such as security and compute power while ensuring that internal datasets are high quality and available for training and use
- Taking an ethical approach and ensuring that bots align with organizational goals and values, as well as involving employees to gain their buy-in
- Starting small, with pilot projects that have compelling use cases, and then growing to develop a range of T-shaped bots focused on specific subject areas

Overall, integration of AI-powered knowledge bots within organizations holds immense potential to transform knowledge management, empower employees, and drive organizational success in both the long and short term. Finally, the knowledge management challenge that organizations have been grappling with for decades is on the verge of being solved.

NICHOLAS JOHNSON

is a Partner in Arthur D. Little's London office and a member of the Digital Problem Solving (DPS) Practice.

MICHAEL PAPADOPOULOS

is the Chief Architect in Arthur D. Little's London office and a member of the Digital Problem Solving (DPS) Practice.

FOIVOS CHRISTOULAKIS

is a Consultant in Arthur D. Little's London office and a member of the Digital Problem Solving (DPS) Practice.

GREGORY BROWN

is a Business Analyst in Arthur D. Little's London office and a member of the Digital Problem Solving (DPS) Practice.

ANASTASSIYA TIKHONOVA

is a Business Analyst in Arthur D. Little's London office and a member of the Digital Problem Solving (DPS) Practice.

EYSTEIN THANISCH

is a Business Analyst in Arthur D. Little's London office and a member of the Digital Problem Solving (DPS) Practice.