Transforming Knowledge Management in Business with Generative AI: Alternative Solutions, Risks and Regulatory Considerations

Published: 05.09.2024

Umair Ali Khan, Senior Researcher, Haaga-Helia University of Applied Sciences

Dmitry Kudryavtsev, Senior Researcher, Haaga-Helia University of Applied Sciences

In our previous article, based on insights from the seminar on "Generative AI-Enhanced Knowledge Management in Business" in Haaga-Helia University, we explored the opportunities that generative AI presents for business knowledge management. We discussed practical use cases and the challenges businesses face when transitioning to generative AI systems, particularly the complexities of change management and system integration.

In this article, we build on those insights by examining alternative technological solutions for implementing generative AI, the risks and regulatory considerations involved, and the critical success factors that can drive successful AI adoption in businesses. By understanding these aspects, businesses can better navigate the challenges of generative AI implementation and harness its full potential.

Alternative Technological Solutions to Generative AI Implementation in Business

Various solution types can be used to implement Generative AI into business: Retrieval Augmented Generation (RAG), LLM fine-tuning, Knowledge Graphs, and combinations of the aforementioned approaches.

RAG ensures that the AI solution is aware of company-specific and use case specific data or knowledge (Herman, 2023). Incorporating this data directly into the LLM is often impractical due to high costs and the inherent risk of inaccuracies in LLM-generated responses. LLMs operate based on statistical models, which means there is always a possibility of generating incorrect answers. By integrating RAG, businesses can ensure that the LLM has access to accurate, up-to-date information relevant to their specific needs.

Fine-tuning a custom LLM is crucial for improving how well the model follows instructions and ensuring it communicates in the company's preferred language and tone of voice. Fine-tuning addresses specific demands, particularly instruction following, which is essential for LLMs to perform effectively in narrow, specialized use cases.

The solutions of **LLMs with fine-tuning** (Raj J, VM, Warrier, & Gupta, 2024) and **RAG applications** typically **coexist** because they serve complementary functions. In practical terms, even with a RAG system, an LLM is required to formulate responses to queries. By default, an LLM can perform adequately, but specific use cases may require tailored outputs. For instance, in customer support, responses need to be short, concise, and to the point. Conversely, in marketing, the generated text should be longer, more engaging, and sales-oriented. While some of these adjustments can be managed through prompt engineering, fine-tuning the LLM is essential for achieving a good performance. Fine-tuning allows for precise control over the LLM's outputs, ensuring they meet the specific requirements of different business applications. Moreover, **combining a fine-tuned custom LLM with a RAG system** can offer an alternative solution. This approach leverages the strengths of both methodologies: the specificity and relevance provided by RAG and the refined, context-appropriate responses facilitated by fine-tuning. Together, they create a robust and adaptable AI solution tailored to a company's unique needs and objectives.

Another promising solution which addresses the limitations of both RAG and fine-tuned custom LLMs is the **integration of knowledge graphs into LLMs** (Zhu et al., 2023). Knowledge graphs provide a precise way of retrieving relationships between pieces of data. This means that businesses can combine general knowledge of LLMs with exact knowledge and relationship of the company's business data. This process includes indexing the data to create sub-summaries of various topic clusters and establishing links between these clusters and the organizational structure, detailing who is working on which topics and what data sources they use daily. By doing this, the system can utilize user information to rank context more effectively, prioritizing the most relevant and accurate information. Although this integration requires significant effort, it is a natural progression for enhancing the accuracy and factuality of the data. The resulting knowledge graph-based system provides a more detailed, interconnected view of the company's data, leading to improved information retrieval, more precise insights, and ultimately, more informed decision-making processes.

Another dimension of technological decisions is associated with the selection of LLM. While only a few companies have the capability to build generative AI models from scratch, the availability of open-source LLMs provides a viable starting point. Businesses can begin with these open-source models and then build upon them, tailoring the technology to meet their unique needs. This approach allows companies to develop customized AI solutions that align with their strategic goals, making the process more manageable and beneficial for gaining a competitive advantage.

Some companies cannot use public clouds due to security considerations or strategic reasons. Therefore, they should not solely rely on major providers like Microsoft or Amazon for AI infrastructure. When companies depend on a single provider for LLMs or cloud platforms, they follow the provider's roadmap, not their own. To truly lead in the AI space, companies need to own their Generative AI technology and develop the capabilities to control and innovate within their unique frameworks. This approach ensures that companies remain independent and strategically agile, allowing them to tailor their AI implementations to their specific needs and objectives.

Concerning the future of transparency in AI models, particularly the **shift toward open-source models**, some companies rely on open-source models due to the inability to use proprietary models on-premises and within their products. This reliance underscores a broader industry discussion about what constitutes "open source," given that not all model weights or data sets for open-source models like Llama 3 are fully disclosed.

Open-source models will prevail in the long run, primarily because they offer the necessary transparency that proprietary models lack. Transparency is crucial for building trust and ensuring the integrity of AI systems. Open-source models allow for greater scrutiny and understanding of how the models operate, which is essential for mitigating risks and enhancing accountability.

The historical success of open-source approaches in various technological domains supports this perspective. By enabling more collaborative and transparent development processes, open-source models can better address the ethical and operational challenges associated with AI. As the demand for transparency continues to grow, the industry will increasingly favor open-source solutions, ensuring a more open and reliable AI landscape in the future.

GenAI-related Risks and Regulatory Complications

One of the primary concerns with generative AI is the phenomenon of "**hallucinations**," where AI generates information that appears plausible but is incorrect or misleading. This unreliability poses a significant problem, necessitating the development of more **trustworthy AI systems** and raises the need for external regulation.

The question of regulation is complex and multifaceted, involving various jurisdictions and legal frameworks. It raises concerns about whether regulation should be handled at the European Union level, under EU law, or under American or Chinese rules. The regulatory landscape is a "big, confusing pot" of different rules and standards, making it challenging for companies and users to navigate.

For the average user and businesses, a practical approach is grounded in common sense, where inputting any business-critical information into proprietary LLMs is not recommended. Essentially, users should avoid sharing anything with AI that they wouldn't comfortably share with a stranger. This cautious approach helps mitigate risks while leveraging the benefits of generative AI.

The **new ethical regulations, such as the AI Act** (European Parliament, 2024), on LLM-based solutions are welcoming, particularly for sensitive markets like finance and healthcare. The regulations mandate traceability for decisions made by LLMs, ensuring that there is a clear record of how each decision was formed. This requirement helps in managing hallucinations and monitoring the performance of AI models, ensuring their outputs are accurate and reliable.

Implementing various **fact-checking and guardrail systems to ensure the integrity of AI-generated content** is important. For large enterprises, data security is a paramount concern. When integrating data from multiple sources (e.g., in a RAG application), it is crucial to manage who can access what information. Traditional databases often lack user management or permissions based on individual

users. Therefore, it is necessary to establish rules that control data access dynamically, ensuring that sensitive information is protected.

Regulations in creating a secure environment for AI applications are of particular importance. Ensuring that information does not leak through LLM apps is a critical aspect of these regulations. Such regulatory measures would become mandatory in the long term, providing a structured framework to safeguard data security and maintain the integrity of AI operations. This proactive approach to regulation can help mitigate risks while enabling the benefits of generative AI.

Bridging the Generative AI Readiness Gap

The evolution of AI readiness in organizations has seen significant changes, especially with the advent of generative AI technologies. Initially, discussions around AI readiness were more prevalent. However, as generative AI became more accessible (e.g. via ChatGPT interface and APIs), many companies began initiating proof of concepts (POCs) without revisiting these crucial readiness components.

In terms of critical requirements for GenAI adoption, it is essential to go beyond basic data readiness. Additional competencies, such as knowledge modeling, are becoming increasingly important as companies aim to advance their AI implementations. There is a notable gap in AI readiness among organizations, particularly for generative AI. While tools like GPT have prompted many companies to start POCs, **there is often a lack of comprehensive understanding of the broader implications and requirements of these technologies**.

This rush to adopt AI has resulted in varied levels of preparedness across organizations. **Many started POCs without fully grasping what it takes to effectively integrate AI into their operations, leading to disparities in AI readiness. However, organizations are rapidly catching up, especially in areas impacting their business operations**. Conversely, projects aimed at addressing long-term gains face more significant challenges. Since these projects often do not align with immediate business priorities insufficient effort and resources may hinder the growth of needed AI readiness.

Data readiness remains a critical factor for traditional data science and machine learning tasks. However, companies produce relevant data for generative AI on a daily basis and often have the necessary data readily available for such AI applications. **In the case of on-premise implementation of generative AI solutions, the use of modern infrastructure solutions, like virtualization technologies such as Kubernetes, is crucial for effectively deploying GenAI.** Companies vary widely in their adoption of these technologies. Those lacking modern infrastructure solutions may find it challenging to create effective on-premise Generative AI solutions.

Critical Success Factors of Generative AI Implementation

The successful implementation of AI hinges on several critical factors:

- **Focusing on high-quality, well-managed data is fundamental,** as it forms the foundation of effective AI applications.

- **Aligning AI initiatives with clear business cases that address genuine customer needs** ensures that these technologies deliver real value.
- **Strategic thinking and strong managerial support are also essential**, as they provide the necessary framework and leadership for systematic adoption.
- **A long-term perspective, coupled with well-defined KPIs**, is vital to tracking progress and sustaining success.

Conclusion

While generative AI holds immense potential, successful adoption requires careful navigation of the underlying challenges. Businesses must select the right technological approach, whether it's retrieval augmented generation, fine-tuning of language models, or the integration of knowledge graphs, to suit their specific needs. Additionally, understanding and mitigating risks, particularly those related to AI "hallucinations" and data security, is essential to maintaining trust and ensuring compliance with emerging regulations.

Ultimately, the key to unlocking the full potential of generative AI lies in aligning its implementation with clear business objectives, supported by strong governance and a culture of continuous learning and adaptation. By taking a strategic and thoughtful approach, businesses can leverage the transformative power of generative AI for driving innovation and maintaining a competitive edge in the evolving landscape of knowledge management.

References

European Parliament 2024, March 8. Artificial Intelligence Act: MEPs adopt landmark law. Retrieved May 29, 2024.

Herman, J. 2023. Retrieval-Augmented Generation (RAG): From theory to LangChain implementation. Towards Data Science. Retrieved May 29, 2024.

Raj J, M., VM, K., Warrier, H. & Gupta, Y. 2024. Fine tuning LLM for enterprise: Practical guidelines and recommendations. arXiv preprint.

Zhu, Y., et al. 2023. LLMs for knowledge graph construction and reasoning: Recent capabilities and future opportunities. arXiv preprint.