


**Whitepaper**

# Harnessing Large Language Models (LLMs) for Decision Intelligence Advancements



# Contents

---

1. Abstract	03
2. Introduction	04
3. Large language models and Decision Intelligence: The future of business decision making	06
4. Challenges in incorporating LLMs into DI Applications	08
5. LLM-Powered Applications: Frameworks, Strategies, and Project Lifecycle	10
A. Frameworks for LLM-Based applications and benchmarking strategies	10
B. Build vs. Buy Dilemma: Navigating LLM development and integration	12
C. Navigating the Generative AI Project Lifecycle	14
6. Melding Generative AI with Decision intelligence - How Lumin utilized LLMs to supercharge business outcomes	14
7. Conclusion	16
8. References	16
9. About the author	17

# Executive Summary

---

The surge in interest in generative AI (Gen AI) is primarily due to the rapid adoption of large language models (LLMs), such as GPT-4. There has been intense discussion among experts regarding these models; some view them as the forthcoming revolution in artificial intelligence (AI), while others emphasize their potential risks and constraints. In an era of data-driven decision-making, the fusion of artificial intelligence (AI) and natural language processing (NLP) has become a potent catalyst for progress. LLMs have emerged as a pivotal asset in this paradigm shift, promising to augment decision intelligence (DI) applications in unprecedented ways.

This potent combination comes with its own set of challenges. This whitepaper serves as a guide for organizations and decision-makers, offering insights into the multi-faceted world of LLMs and DI.

Whether you are considering the build vs buy dilemma, evaluating LLM frameworks, or seeking to understand the Gen AI project lifecycle, this whitepaper equips you with the knowledge needed to navigate the evolving landscape.



# Introduction

---

LLMs have ushered in a new era of decision-making by enhancing natural language understanding (NLU), predictive analytics, and contextual reasoning. Innovative organizations are already using decision intelligence to improve their speed to market and enhance their decision-making. Large Language Models (LLMs) represent a major trend in this field due to their ability to harness massive datasets of text, ML models, and advanced analytics to perform diverse tasks such as text generation, translation, and summarization, making it faster to enable data-driven decisions.

This integration amplifies the potential for uncovering valuable insights, thus augmenting the impact and effectiveness of decision-making processes across a broad spectrum of industries and domains. However, the convergence of LLMs and DI is not without its complexities. While LLMs offer immense potential, they also bring forth a host of challenges, ranging from computational resources, model evaluation, interpretability, and ethical considerations. Given the constraints, organizations willing to experiment with this technology must understand its complexities and implications and assess their ability to embark on a full-fledged generative AI project.

This whitepaper provides a comprehensive understanding of how LLMs can revolutionize decision intelligence, the challenges and considerations involved, and a roadmap for successfully integrating LLMs into DI projects.

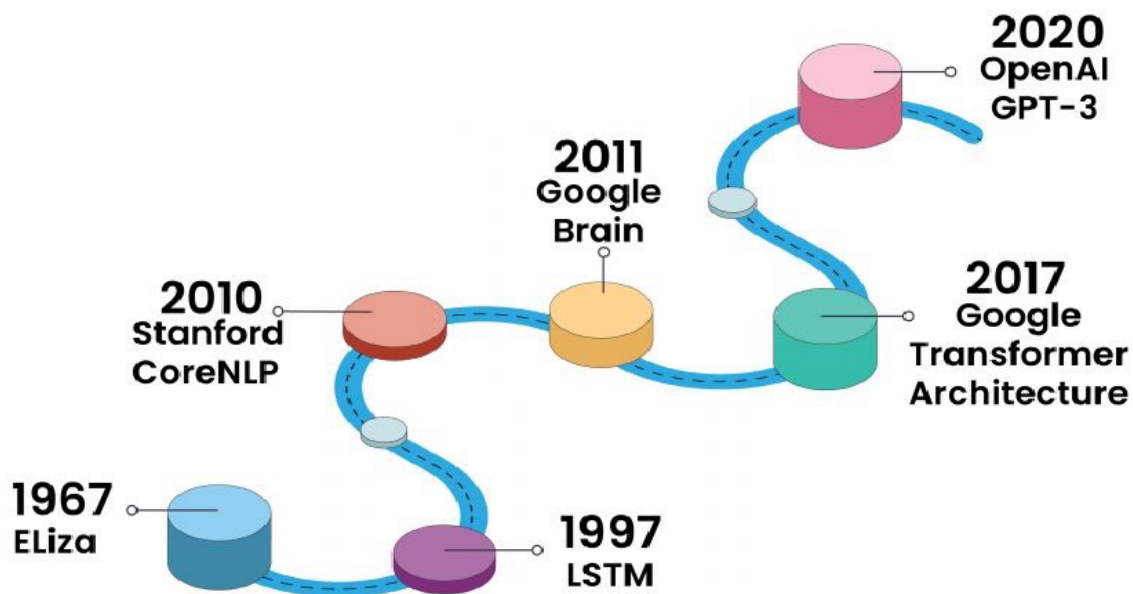
## What are Large Language Models (LLMs)?

---

Large language models (LLMs) represent a category of artificial intelligence (AI) systems that have undergone training on extensive volumes of textual data. . These models are also called neural networks (NNs), which are computing systems inspired by the human brain. These neural networks use a network of layered nodes, much like neurons. They use transformer models and are trained using massive datasets. i This capability allows them to identify, translate, forecast, or produce text and various forms of content.LLMs employ sophisticated machine learning (ML) algorithms to comprehend and scrutinize the subtleties within human speech, encompassing syntax, semantics, and contextual meaning.

## A brief history of LLMs

The history of Large Language Models can be traced back to 1960 when AI and natural language processing experiments were rudimentary. There has been notable advancement in the realm of AI and neural networks over time. In 2018, OpenAI unveiled GPT-1, an important development in natural language processing (NLP) capabilities. GPT-2 followed, demonstrating the capacity to generate coherent and sometimes convincing text even when it contained errors or biased language. Finally, in 2020, GPT-3 was released, which proved to be a game-changer for NLP. GPT-3's vast parameter set enables it to perform an impressive range of tasks, from language translation and text summarization to question answering and creative writing.<sup>ii</sup>



## A brief history of LLMs

LLM already has a variety of real-world applications, from supporting customer service and improving operational services to generating personalized content or accelerating product development. As this technology progresses further, it will radically transform how we communicate, generate content, consume information, and even comprehend the world around us.


Unlike traditional AI, which recognizes patterns in existing data to make predictions, generative AI, powered by LLMs, goes beyond what already exists to create something new. This is done by training the AI on a vast collection of existing content, then employing the acquired patterns to generate fresh content that mirrors similar styles or themes. GenAI finds applications in tackling intricate issues across diverse fields, including drug discovery, climate modeling, and engineering. As generative AI technology continues to progress, we can anticipate witnessing further inventive applications of this technology in the future.

# Large language models and decision intelligence —The future of business decision-making

In the ever-evolving business landscape, the speed and efficiency of decision-making are constantly challenged. Conventional approaches such as business intelligence (BI), reliant on data, analytical models, and expensive resources, tailored for the existing state, are no longer satisfactory. With advancements in technology, specifically LLMs, the potential for AI and machine learning to enhance decision-making is astounding. LLMs, exemplified by powerful Generative AI constructs (like GPT-3.5), represent the pinnacle of natural language understanding and generation. They possess the unique ability to decipher human language intricacies, context, and intent at an unprecedented scale. Concurrently, decision intelligence, the art of leveraging data and analytics for decision-making, is the cornerstone upon which business strategies are built.

Until recently, LLMs were just prediction machines. But, they are now transforming how AI tools interact with humans, considering the nuances and context of business communication, especially related to organizational decision-making processes. By faster data processing and analysis, LLMs can provide decision-makers with accurate, timely, and context-aware insights, ultimately leading to better decisions.

Generative AI, powered by large language models, can enhance decision intelligence systems in several ways. Here are a few examples:



## Augment AI training data sets

Gen AI can be used to generate new data points that can be used to improve the accuracy and reliability of DI models. This proves especially advantageous in cases where actual data is limited or challenging to acquire. Additionally, it has the capacity to simulate diverse scenarios and produce extensive data sets, enhancing the training of AI models for more insightful decision-making.

## Enhance contextual learning

To be effective, any decision intelligence tool needs to understand user and business context to give meaningful responses. Gen AI models can also be used to understand user questions better. It can provide more accurate results back to the users of decision intelligence tools. The complex AI model responses or recommendations generated by these tools can also be simplified using Gen AI technologies to understand and interpret their reasoning.

## Enrich text analysis

One of the most effective uses of LLMs is text analysis. LLMs are particularly adept at quickly processing large amounts of text and extracting key insights. This is achieved through natural language processing techniques, which allow the algorithm to understand language patterns and interpret the meaning. When combined with decision intelligence, text analysis can be incredibly useful in various industries, from finance to healthcare. For example, LLMs can extract insights from medical records to help doctors diagnose better.

## Improve interactivity

Gen AI can generate more human-like responses created by NLP models that can be used in DI systems. These responses make it easier for users to interact and communicate with the system and get relevant information. Natural language generation (NLG) involves creating written or spoken language using structured data as its basis. For example, an LLM can be trained to generate a report summarizing financial data in natural language. For businesses relying on financial statements and reports, this capability can result in notable savings of time and resources. LLMs can also generate customer service responses, news articles, automatically generate reports and charts based on the user's data, thereby providing a fast and efficient way to communicate insights and results.

Integrating LLMs with decision intelligence represents a pivotal leap forward in AI-powered decision-making. It addresses the prior limitations and opens doors to a myriad of possibilities in data-driven decision-making. By fusing the power of LLMs with the power of NLG and ML-powered intuitive interfaces, business users can now access a comprehensive and flexible toolset that adapts to their evolving needs.

# Challenges in incorporating LLMs into DI applications

Incorporating LLMs into DI applications to enhance its capabilities is inherently a generative AI problem. This endeavor presents unique challenges and promising opportunities within the domain of decision intelligence. The challenge lies in harnessing the full potential of LLMs to enable dynamic and contextually aware conversations while ensuring seamless integration with the DI application's existing architecture. This necessitates a structured approach that navigates the complexities of deploying advanced AI models.

## Domain challenges

### In-depth knowledge of LLM frameworks

The successful integration of LLMs into applications requires robust frameworks and benchmarking strategies. These tools facilitate seamless deployment, and help gauge the effectiveness and efficiency of LLM-based systems. However, scarce knowledge and lack of expert guidance make implementing it in organizations difficult.

### Build vs. buy dilemma

Deciding whether to build an LLM from the ground up or fine-tune an existing model is pivotal. This choice involves critical considerations, including assessing the required expertise and team resources for seamless LLM development and integration.



## Data Challenges

### Computational resources

LLMs demand substantial computational resources for both training and inference processes. . It requires substantial investment in infrastructure, compute resources, and experts to handle and train advanced ML and neural network models.

### Fine-tuning and adaptation

Adapting LLMs to particular tasks or domains can pose challenges, demanding expertise and meticulous consideration to ensure the model performs as intended for the desired application

### Bias and fairness

LLMs can mirror inherit biases in their training data. It's crucial to address these biases and ensure fairness in decision-making, particularly in applications that have significant impacts on individuals or communities.

### Data privacy and security

Using LLMs often involves handling sensitive data. Ensuring data privacy and security throughout the lifecycle of the model is critical and can be a complex undertaking.

### Interpretability and explainability

Comprehending and elucidating the decisions produced by LLMs is intricate. Often categorized as "black box" models, many LLMs present challenges in clearly explaining their outputs.

### Hallucination

Despite their impressive capabilities, LLMs introduce challenges when incorporated into DI applications. One notable hurdle is the potential for hallucinations. These models may generate diverse and sometimes unexpected outputs that may not align with the intended context. Addressing this issue requires a nuanced approach to fine-tuning and careful consideration of the training data to ensure accuracy and reliability in generating desired outcomes.

Developing an LLM-based application entails surmounting several challenges. Selecting the appropriate framework and establishing precise evaluation metrics and benchmarks is imperative to ensure optimal application development and performance.

# LLM-powered applications Frameworks, strategies, and project lifecycle

This section explores two prominent LLM frameworks: industry benchmarks and evaluation criteria. This will help us understand how organizations should address the build vs. buy dilemma regarding LLM projects and the various steps in a typical generative AI project journey.

## LLM frameworks

LangChain and LlamaIndex are two popular LLM frameworks that offer unique capabilities to enhance the utilization of LLMs in applications.

### **LangChain—Empowering Context-Aware Reasoning**

LangChain is a dynamic framework tailored for crafting applications driven by language models.<sup>iv</sup> It empowers applications to be context-aware, establishing connections between a language model and contextual sources such as prompt instructions, exemplar inputs, or other content for grounded responses. Additionally, LangChain enables reasoning, allowing the language model to deduce optimal answers based on the context and make informed decisions. A LangChain features modular components that offer a range of implementations for diverse applications and off-the-shelf chains, providing structured frameworks for accomplishing specific higher-level tasks. This combination simplifies both the initiation of projects and the customization of intricate applications, reinforcing the adaptability of LangChain.<sup>3</sup>

### **LlamaIndex—Structured Data Access for LLM Applications**

Formerly known as GPT Index, LlamaIndex is a pivotal data framework designed to facilitate the ingestion, organization, and retrieval of private or domain-specific data in LLM applications. LLMs often necessitate the integration of private or specialized data sources, which may be dispersed across various applications and data repositories. These may include APIs, SQL databases, or even encapsulated within PDFs and presentations. LlamaIndex steps in to streamline this process by providing data connectors to ingest diverse data sources, data indexes to structure information for optimized LLM consumption, and engines for intuitive natural language access to this data. Moreover, LlamaIndex harmoniously integrates with other elements of the ecosystem, be it LangChain, Flask, Docker, ChatGPT, or any compatible framework, ensuring seamless collaboration within the broader application landscape.<sup>4</sup>

## Evaluating LLMs and benchmarking strategies

Assessing the performance of large language models is critical for ensuring their effectiveness in real-world applications. Two widely recognized evaluation metrics, Rouge and Bleu scores, provide quantitative measures of the quality of the generated text.

### Rouge and Bleu Scores—Quantifying Language Model Performance

Rouge (Recall-Oriented Understudy for Gisting Evaluation) evaluates the quality of generated text by comparing it against human-written reference texts.<sup>iii</sup> It focuses on aspects like overlap in n-grams and word sequences, providing a quantitative assessment of the similarity between generated and desired outputs. Alternatively, Bleu (Bilingual Evaluation Understudy) gauges the quality of machine-translated text. It quantifies the proximity between the generated text and one or multiple reference translations. . Rouge and Bleu scores are valuable tools for gauging the precision and fluency of LLM-generated text, enabling iterative improvements in model performance.

### Evaluation Benchmarks—Beyond Individual Metrics

Standardized benchmarks are essential for a comprehensive evaluation of LLMs.

- The General Language Understanding Evaluation (GLUE) and its advanced counterpart, SuperGLUE, assess a model's proficiency in various language tasks, enabling fair comparisons between different LLM models.
- The Multimodal Multilingual Language Understanding (MMLU) benchmark expands evaluation to include multimodal tasks, considering both text and images.
- Additionally, the Harnessing Evaluation of Large Models (HELM) benchmark thoroughly assesses LLMs across diverse tasks, providing nuanced insights into their capabilities.

When combined with Rouge and Bleu scores, these benchmarks offer a holistic understanding of LLM performance. The understanding aids researchers and developers in making informed decisions regarding model selection and optimization for specific applications. It ultimately drives advancements in natural language understanding and generation. The figure below shows the latest leaderboard from the SuperGLUE benchmark.

Leaderboard Version: 2.0

Rank	Name	Model	URL	Score	BoolQ	CB	COPA	MultiRC	ReCoRD	RTE	WIC	WSC	AX-b	AX-g	
1	JDExplore d-team	Vega v2		91.3	90.5	98.6/99.2	99.4	88.2/62.4	94.4/93.9	96.0	77.4	98.6	-0.4	100.0/50.0	
+	2	Liam Fedus	ST-MoE-32B		91.2	92.4	96.9/98.0	99.2	89.6/65.8	95.1/94.4	93.5	77.7	96.6	72.3	96.1/94.1
3	Microsoft Alexander v-team	Turing NLR v5		90.9	92.0	95.9/97.6	98.2	88.4/63.0	96.4/95.9	94.1	77.1	97.3	67.8	93.3/95.5	
4	ERNIE Team - Baidu	ERNIE 3.0		90.6	91.0	98.6/99.2	97.4	88.6/63.2	94.7/94.2	92.6	77.4	97.3	68.6	92.7/94.7	
5	Yi Tay	PaLM 540B		90.4	91.9	94.4/96.0	99.0	88.7/63.6	94.2/93.3	94.1	77.4	95.9	72.9	95.5/90.4	
+	6	Zirui Wang	T5 + UDG, Single Model (Google Brain)		90.4	91.4	95.8/97.6	98.0	88.3/63.0	94.2/93.5	93.0	77.9	96.6	69.1	92.7/91.9
+	7	DeBERTa Team - Microsoft	DeBERTa / TuringNLRv4		90.3	90.4	95.7/97.6	98.4	88.2/63.7	94.5/94.1	93.2	77.5	95.9	66.7	93.3/93.8
8	SuperGLUE Human Baselines	SuperGLUE Human Baselines		89.8	89.0	95.8/98.9	100.0	81.8/51.9	91.7/91.3	93.6	80.0	100.0	76.6	99.3/99.7	
+	9	T5 Team - Google	T5		89.3	91.2	93.9/96.8	94.8	88.1/63.3	94.1/93.4	92.5	76.9	93.8	65.6	92.7/91.9
10	SPoT Team - Google	Frozen T5 1.1 + SPoT		89.2	91.1	95.8/97.6	95.6	87.9/61.9	93.3/92.4	92.9	75.8	93.8	66.9	83.1/82.6	

Fig 2. Leaderboard on the SuperGLUE benchmark. Source Image from <https://super.gluebenchmark.com/leaderboard>

## Build vs. buy: Navigating LLM development and integration

Incorporating large language models into decision intelligence applications marks a pivotal leap forward in AI-powered decision-making. However, the question is whether an organization should invest in building LLM capabilities in-house or leverage existing solutions like GPT or any other foundation models. This decision needs careful consideration and assessment of existing resources and capabilities.

### a) Expertise and team requirements

Building LLMs necessitates a multidisciplinary team with expertise in several key domains:

- Machine learning engineers: They are proficient in designing and training deep learning models, particularly LLMs, to meet specific application requirements.
- Natural Language Processing (NLP) specialists: They fine-tune and adapt LLMs to domain-specific contexts, ensuring accurate and contextually relevant responses.
- Data scientists: They are skilled at preprocessing and structuring data for model training and evaluating model performance using relevant metrics.
- Domain experts: Their domain-specific knowledge is crucial for crafting effective prompts and instructions to guide LLM behavior accurately.
- Ethical AI specialists: They address bias, fairness, and ethical considerations in LLM deployment, ensuring responsible AI practices.
- Software developers: They are required to integrate LLMs into existing DI applications, optimizing performance and ensuring scalability.
- Security experts: They play a crucial role in protecting sensitive data, establishing resilient data privacy protocols, and guaranteeing adherence to pertinent regulations.
- UX/UI designers: They design user interfaces for seamless interaction with LLM-powered features within DI applications.

b) Building in-house vs. utilizing existing solutions- The decision to build LLM capabilities in-house or leverage existing solutions depends on various factors.

<b>Consideration</b>	<b>Building In-House</b>	<b>Utilizing Existing Solutions</b>
<b>Expertise and Resources</b>	Requires hiring and training of a specialized team	Requires less in-house expertise
	Provides complete control over model development	Customization options may be limited
<b>Cost Considerations</b>	Significant upfront costs for hiring, training, and infrastructure	May involve licensing fees, with potentially lower initial costs
	The long-term expenses can fluctuate depending on maintenance and support requirements	Licensing agreements and usage may influence costs
<b>Time to Market</b>	Development timelines may be extended due to the learning curve and iterative process	Allows for faster deployment, potentially accelerating time-to-market
<b>Customization and Control</b>	Offers complete customization and control over the LLM	Provides a ready-to-use solution, with customization options potentially constrained
<b>Maintenance and Support</b>	Requires ongoing maintenance, updates, and support for the developed LLM	The solution provider typically provides maintenance and support
<b>Integration Complexity</b>	Integration may be smoother due to intimate knowledge of the organization's tech stack	Integration may require adaptation to fit within existing systems

Table 1. Key considerations and comparison for building vs. buying existing LLM-infused applications

## How to get started—Navigating the generative AI project lifecycle

The generative AI project life cycle offers a systematic and strategic approach to integrate large language models into decision intelligence applications seamlessly. This structured process involves careful planning, robust data governance, and a deep understanding of the underlying technology and the specific application domain. By following this comprehensive strategy, organizations can effectively address challenges related to LLM integration while maximizing opportunities. This includes ensuring data privacy, mitigating biases, optimizing computational resources, and fine-tuning models for specific tasks or domains. Furthermore, it encourages continuous learning and adaptation, allowing LLMs to evolve alongside changing requirements. The four stages of the Gen AI project life cycle serve as a roadmap for harnessing the potential of LLMs in problem-solving scenarios. Figure 3 shows the flow diagram for the Gen AI project life cycle.

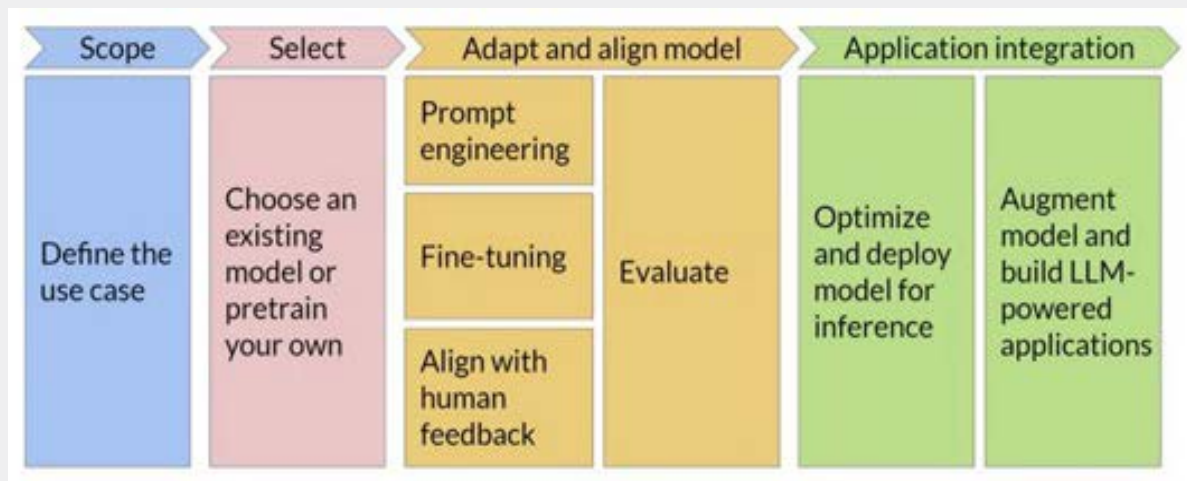


Fig 3. Generative AI Project Life Cycle. Source - Image from DeepLearning.AI

### 1. Scope—Define the use case

At the outset, it is imperative to delineate the precise use case that will drive the incorporation of LLMs into DI applications. This requires a comprehensive grasp of the particular problem domain and the intended objectives. . By clearly defining the scope, stakeholders can establish a focused trajectory for the project, ensuring that the subsequent steps align with the objectives.

### 2. Select—Choose an existing model or pre-train your own

Once the use case is well-defined, the next step involves selecting the appropriate LLM. This decision hinges on model size, training data, and domain relevance. Organizations may utilize an existing, pre-trained LLM or train a custom model tailored to their specific requirements. This stage lays the foundation for the subsequent phases of the project.

### **3. Adapt and align model—Prompt engineering, fine-tuning, align with human feedback, evaluate the model**

This critical stage focuses on refining the selected LLM to align it with the defined use case. It includes prompt engineering, where the input prompts are carefully designed to elicit accurate and relevant responses. Fine-tuning follows, involving iterative adjustments to the model's parameters based on validation data. Human feedback is invaluable in this process, providing the necessary guidance for model optimization. Rigorous evaluation, employing metrics like perplexity and human judgment, ensures the model's performance meets the desired standards.

### **4. Application integration—Optimize and deploy model for inference, augment model, and build LLM-powered applications**

With a refined LLM, the focus shifts towards seamless integration with DI applications. This involves optimizing the model for efficient inference, ensuring it operates in real-time or near-real-time scenarios. Additionally, the model can be augmented with additional features or components to enhance its functionality within the application. Finally, leveraging the LLM-powered capabilities, organizations can develop various applications that enable users to extract insights and make informed decisions with unprecedented ease and depth.

LTIMindtree's AI and Data products division, Fosfor, has emerged as a leader in leveraging Generative AI and LLMs to elevate its existing applications and offer tailored solutions for businesses aiming to harness the potential of generative AI. Fosfor's decision intelligence product, Lumin, crafted to deliver personalized analytics experience tailored for business users, has further expanded its capabilities. This was achieved by integrating LLMs into its proprietary natural language search engine known as Flow. This integration provides a fluid and contextual conversational experience, mirroring human-like interactions, empowering users to unveil actionable AI-driven insights effortlessly.

## **Conclusion**

Integrating large language models into decision intelligence applications marks a significant milestone in AI-powered decision-making. This whitepaper has provided a comprehensive overview of the potential, challenges, and strategies of leveraging LLMs for enhanced business outcomes. Organizations can unlock unprecedented insights and transform raw data into actionable intelligence by seamlessly integrating LLMs with decision intelligence. The generative AI project life cycle serves as a strategic roadmap, guiding organizations through harnessing the full potential of LLMs. As we look ahead, the continued evolution of LLM technology promises even greater innovation and impact across a broad spectrum of industries. Embracing this transformative technology will undoubtedly be a key driver in shaping the future landscape of decision intelligence.

## References

---

- I. How Generative AI is revolutionizing professional services November 6, 2023:  
<https://achworldwide.medium.com/how-generative-ai-is-revolutionizing-professional-services-32f126e96220>
- II. Natural Language Processing (NLP) Breakthroughs: Understanding the Power of Language Models, Ubaid Ghazali, July 24, 2023:  
<https://medium.com/@ubaidghazali80/natural-language-processing-nlp-breakthroughs-understanding-the-power-of-language-models-3cad82ab023>
- III. The Generative AI LLM Project Lifecycle: From Inception to Application, Haytham Cheiouhou, September 1, 2023:  
<https://medium.com/@haytham99cheikhrouhou/the-generative-ai-llm-project-lifecycle-from-inception-to-application-1e974fba5a7b>
- IV. LangChain- Get started: [https://python.langchain.com/docs/get\\_started/introduction](https://python.langchain.com/docs/get_started/introduction)



## About the author

---



### Ankita Asthana

Product Marketing Lead, Fosfor

Known as a creative individual, Ankita is an experienced product marketing lead, passionate about data storytelling. Ankita has over 10 years of experience in building successful marketing strategies that drive revenue growth and increase brand awareness for data analytics and AI products through effective product positioning and messaging, demand generation, and sales enablement.

LTIMindtree is a global technology consulting and digital solutions company that enables enterprises across industries to reimagine business models, accelerate innovation, and maximize growth by harnessing digital technologies. As a digital transformation partner to more than 700 clients, LTIMindtree brings extensive domain and technology expertise to help drive superior competitive differentiation, customer experiences, and business outcomes in a converging world. Powered by 83,000+ talented and entrepreneurial professionals across more than 30 countries, LTIMindtree — a Larsen & Toubro Group company — combines the industry-acclaimed strengths of erstwhile Larsen and Toubro Infotech and Mindtree in solving the most complex business challenges and delivering transformation at scale. For more information, please visit <https://www.ltimindtree.com/>.