expert.ai

Large Language Models: Opportunity, Risk and Paths Forward

In recent years, large language models (LLMs) such as GPT-3, BERT and T5 have demonstrated remarkable capabilities in natural language processing, transforming the way we interact with language. The release of ChatGPT just six months ago created an easy way for people to experience the power of generative AI and LLMs firsthand. Business and technical executives at organizations large and small are working to quickly determine the most beneficial way to leverage this new age of language-driven AI.

However, along with these opportunities, there are also major concerns related to potential biases, misuse of language models for malicious purposes, unapproved disclosure and use of proprietary information and, last but not least, a lack of truthfulness. This "Large Language Models: Opportunity, Risk and Paths Forward" report summarizes survey results from 300 business, technical and academic natural language AI experts. The goal is to explore potential opportunities and risks associated with LLMs and to provide recommendations for a path forward that enterprises can use in their development and deployment of Large Language Models.

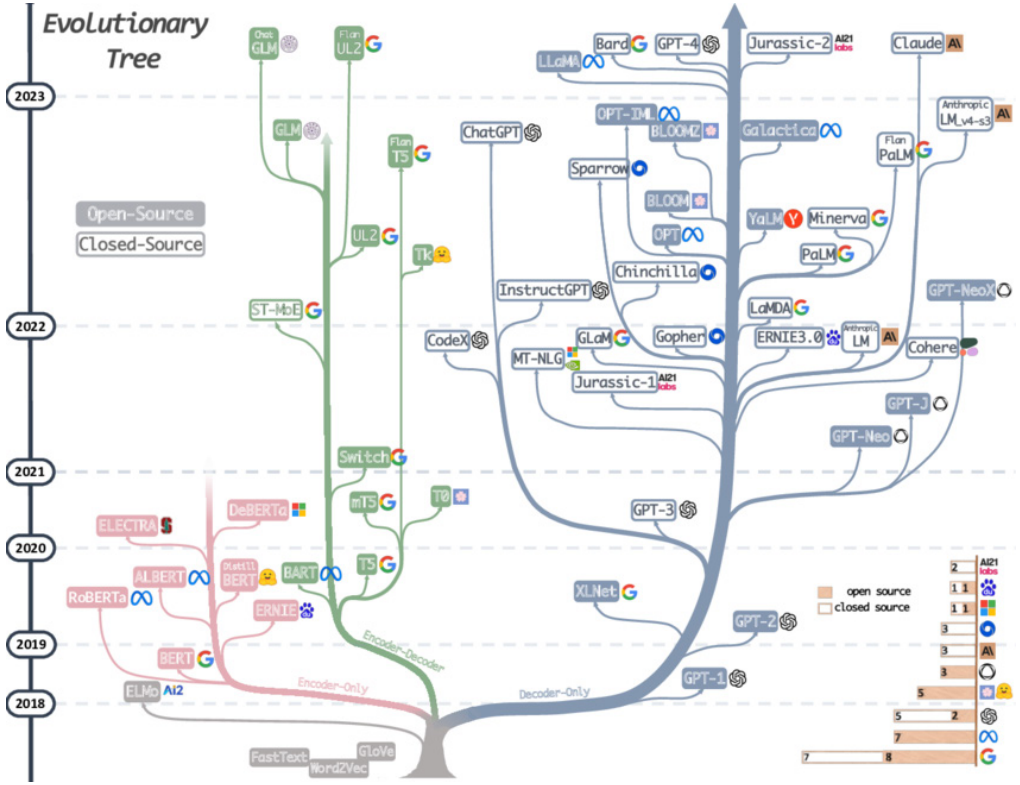

1 Source **"Harnessing the Power of LLMs in Practice: A Survey on ChatGPT and Beyond"**

What are LLM Hallucinations?

In the context of large language models, hallucinations refer to the generation of text or language that is not based on any real-world knowledge or experience. It's how LLMs make things up, or fabricate, in response to prompts or requests.

Few Favor a Training LLM Moratorium but Majority Want Regulation

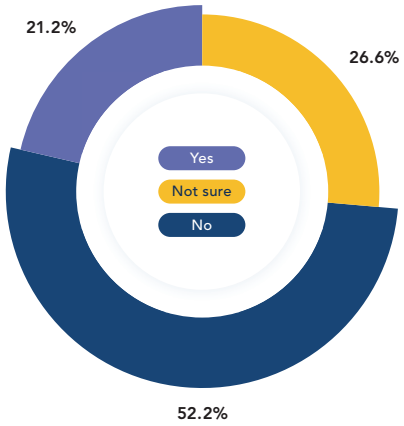

Need for a Training and Development Moratorium

"Should there be a 6-month moratorium on further AI development/training until regulations are in place?"

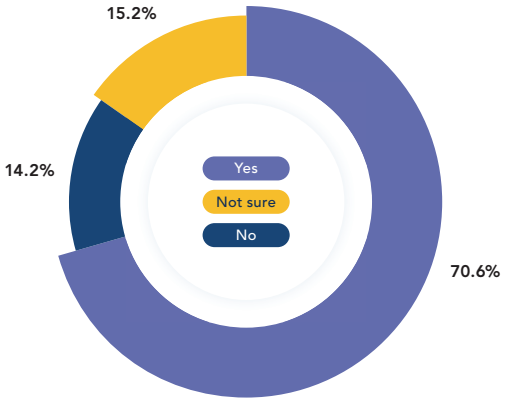

Regulatory Need for Commercial and Malicious Use of AI

"Should government regulations be immediately created to deal with legitimate commercial AI use and malicious use?"

In late March 2023, AI experts, industry leaders, researchers and others signed an **open letter** calling for a six-month "pause" on large-scale AI development beyond OpenAI's GPT-4. Shortly after that, the AI policy group CAIDP (Center for AI and Digital Policy) launched a **complaint with the FTC**, arguing that the use of AI should be "transparent, explainable, fair, and empirically sound while fostering accountability," and that OpenAI's GPT-4 "satisfies none of these requirements" and is "biased, deceptive, and a risk to privacy and public safety." Then, **Italian regulators** called for a block to Chat-GPT access and for OpenAI to address specific privacy and access concerns. OpenAI has subsequently addressed the concerns of Italian regulators and access to Chat-GPT has been restored.

The majority (52.2%) of survey participants stated that there was no need to have a 6-month moratorium on LLM training. Over 70% of AI professionals and practitioners were much more in favor of immediate government regulation to help deal with malicious use and commercial adoption of generative AI. Specific areas of concern for generative AI and LLM adoption are related to truthfulness and veracity of sources, biased output and illegal use of copyrighted materials.

expert.ai

Generative AI and LLMs Can Support Lots of Enterprise Use Cases

Enterprises have a long list of use cases to consider when deciding how to adopt generative AI and LLM capabilities. Typically, they fall into 4 main categories:

Human-Computer Interaction

Providing customers with quick and easy access to information and support, such as product details, troubleshooting guides and frequently asked questions is key to driving the success of digital support channels and CSAT scores. The most prevalent use cases are chatbots (54.4%), question & answering (52.6%) and customer care (22.5%).

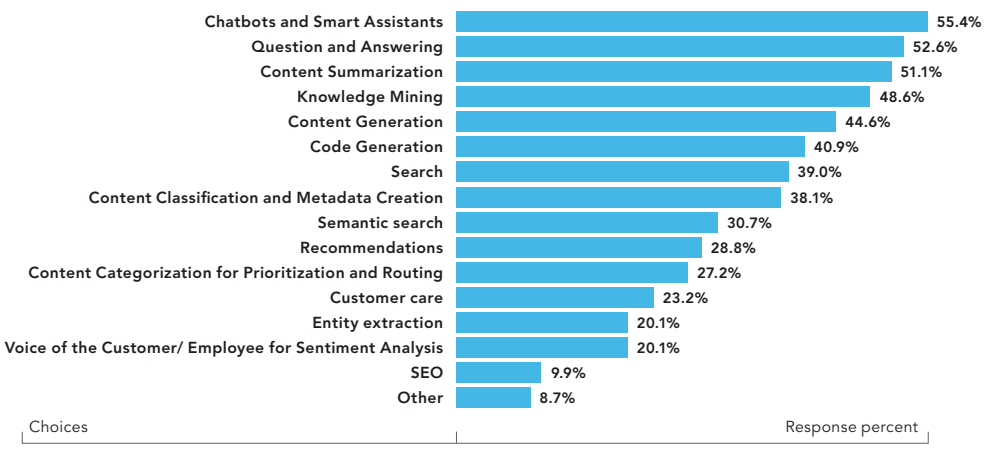

"What are the likely use cases that your company/organization will employ over the next 12-24 months that leverage generative AI and Large Language Models? (Select all that apply)"

Language Generation

Generative AI can write new content, create realistic images, generate marketing copy, compose music and even generate programming code. The two most popular examples of generative AI use cases are content summarization (51.1%) and content generation (44.6%).

Information Extraction

Automatically reading text to identify the topics or entities in a news story, a scientific article or even an email so that they can be sorted into different categories, extracted, routed to the right people or even prioritized based on sentiment and intent are core text analytics use cases. Knowledge mining (48.6%), content classification and metadata creation (38.1%), along with content categorization for routing (27.2%) and entity extraction (20.1%) are all being explored.

Search

Search relies on both content classification and categorization concepts but is typically more task-focused on trying to understand what people are asking for and finding the best resources using natural language queries. General search (39.0%), semantic search (30.7%) and recommendations (28.8%) are important tools for helping people find the information they need quickly and accurately, without having to look through lots of irrelevant results.

Whatever the use case, enterprises still need to make sure that they have a human expert that can verify the correctness of the information.

Significant Adoption Challenges for Generative AI and LLMs Exist

LLMs have been adopted by enterprises for customer service, content creation and data analysis. However, there are several challenges that organizations may face when adopting LLMs at an enterprise level. First and foremost is data privacy and security (73.1%). LLMs typically require a lot of data to be trained effectively, and this data may contain sensitive information. Enterprises must ensure that they have proper data privacy and security measures in place to protect their data from unauthorized access or use.

LLMs are a relatively new technology, and there may be a shortage of skilled professionals (40.7%) who have the expertise to develop and implement them. Enterprises may need to look to outside experts who have experience tuning and placing LLM models and solutions in production (51.2%) and in initial LLM model selection (31.7%). Over the long-term, enterprises that see language as a core asset and differentiator should invest in training or hiring personnel with the necessary skills and knowledge to successfully adopt LLMs.

Enterprise Adoption Challenges for Generative AI and LLMs

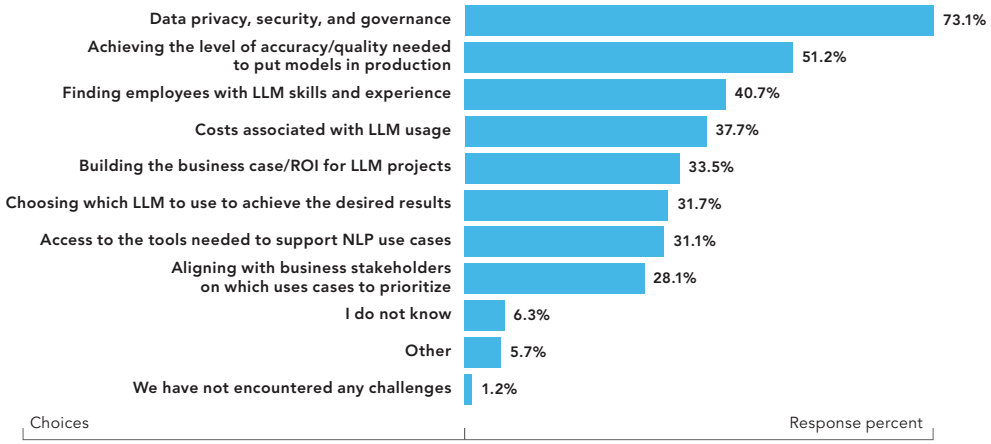

"What are the enterprise adoption challenges you see for generative AI and Large Language Models? (Select all that apply)"

Lastly, LLMs require a lot of computational resources to run (37.7%), which can be expensive and time-consuming to set up and maintain. Enterprises need to have the appropriate infrastructure in place, such as powerful servers or cloud computing services, to support the large-scale deployment of LLMs.

Overall, enterprise adoption of LLMs requires careful planning and consideration for a range of factors, including data privacy and security, infrastructure and resource requirements, integration with existing systems, ethical and legal considerations, and skill and knowledge gaps. To avoid project delays and failed efforts, it makes sense to initially engage with experts in enterprise language model development, tuning and deployments.

expert.ai

ESG Risks Are Real for Generative AI and LLMs

Large language models can have significant ethical and legal implications, particularly around issues of bias, fairness and truthfulness. At the top of the list of concerns among AI professionals was lack of truthfulness (69.8%). Many LLMs like GPT 3.x are trained on wide swaths of information, some of which is copyright protected, and because it comes from publicly available internet data, it has a fundamental *garbage in, garbage out* issue. This copyrighted information was cited as a risk for 59.3% of respondents.

Generative AI models like GPT do not answer questions as much as they guess at the answer that you are looking for, depending on your question. Sometimes, the result can be something humorous or entertaining. Occasionally, it results in false or biased information being presented as fact. And because it's presented in a way that is grammatically correct and authoritative sounding, the end user may not know the difference. Forrester Research calls this '**coherent nonsense**'.

Enterprises must ensure that their use of LLMs is ethical and legal and that they are not inadvertently perpetuating discrimination, copyright infringement or other harmful practices. To mitigate these risks, enterprises need to carefully consider the data used to train LLMs, ensure they have processes in place to identify and address biases, and use additional validation methods to fact check results against known and trusted sources of truth.

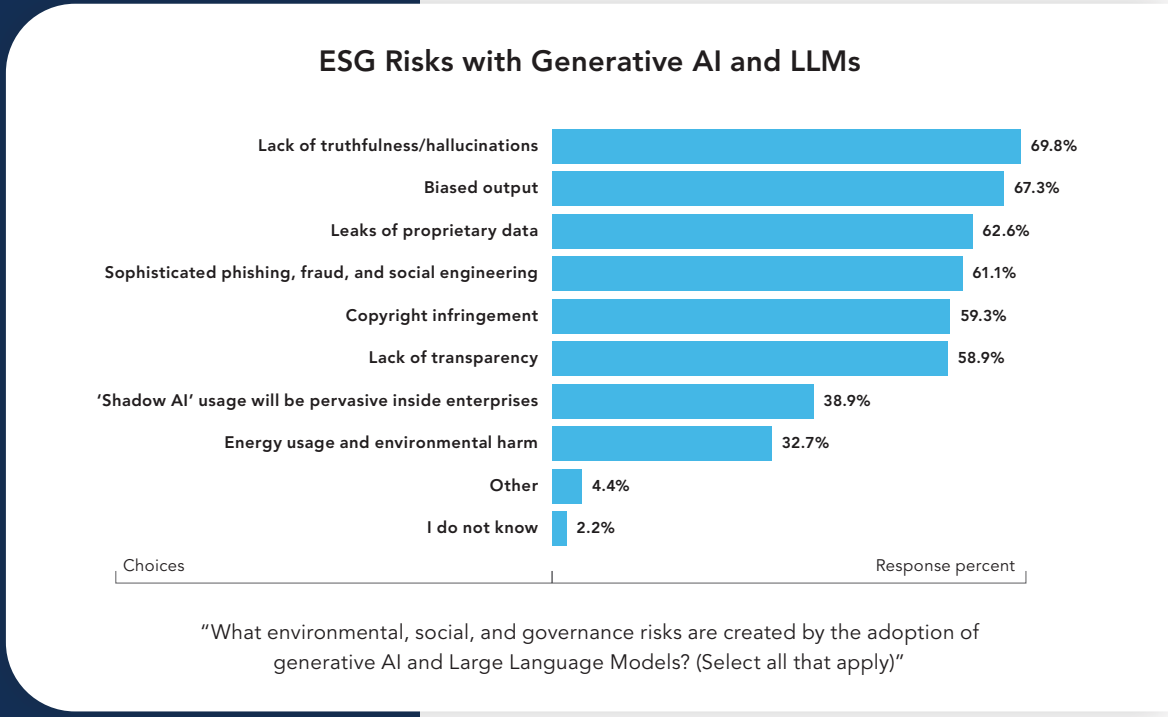

"What environmental, social, and governance risks are created by the adoption of generative AI and Large Language Models? (Select all that apply)"

Governance Principles Still Apply to Generative AI and LLMs

Employees are eager to test drive LLMs to solve problems. Enterprises need to clearly communicate their policies for acceptable use to employees, contractors and partners. Survey respondents fell into 3 groups: 24.0% indicated that further restrictions need to be put in place; 38.8% indicated that some additional degree of freedom should be encouraged; and 34.3% felt that the current principles were adequate or that it was too soon to tell. Regardless of the direction an organization chooses, basic AI data governance principles still apply. Additional governance considerations include:

Model Governance

Enterprises should have clear policies and procedures in place to manage LLMs throughout their lifecycle, including model selection, model validation and model performance monitoring.

Ethical Governance

Enterprises must ensure that the use of LLMs is ethical and aligned with their values and mission. This includes identifying and mitigating potential bias in the data and models, ensuring data privacy and security, and considering the potential social and environmental impacts of LLMs.

Regulatory Governance

Enterprises need to be aware of data protection and privacy regulations, industry-specific regulations and emerging regulatory frameworks related to AI.

Stakeholder Engagement

Enterprises should engage with stakeholders, including employees, customers and partners, to ensure that their use of LLMs is transparent and understandable.

Human Centered

Having humans at only the beginning or only the end of an AI process is not enough to ensure accuracy, transparency or accountability. Enterprises benefit from a human-centered approach, where data and inputs can be monitored and refined by subject matter experts throughout the process.

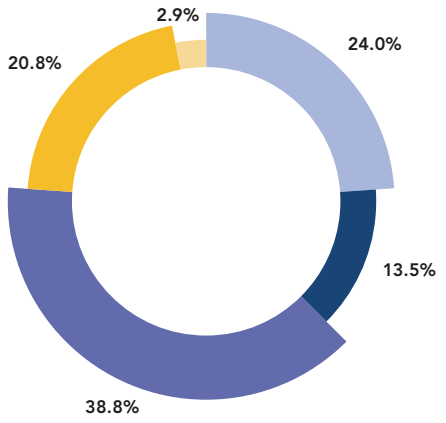

Governance Approach with Generative AI and LLMs

- 24.0% Yes, companies need to restrict teams from using generative AI and Large Language Models services without approval.
- 13.5% No, the same principles apply.
- 38.8% Yes, companies need to encourage teams to explore the potential of generative AI and Large Language Models.
- 20.8% It is too soon to tell.
- 2.9% I don't know.

"Does generative AI and Large Language Models represent a significant change to the governance principles of managing availability, usability, integrity, and security of enterprise data?"

Enterprise-Specific Language Models are the Future

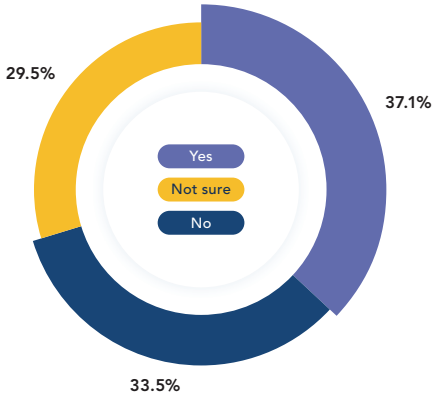

In-House Adoption of Generative AI and LLMs

"Is your company considering building its own business specific Language Model, or adapting/tuning an existing one?"

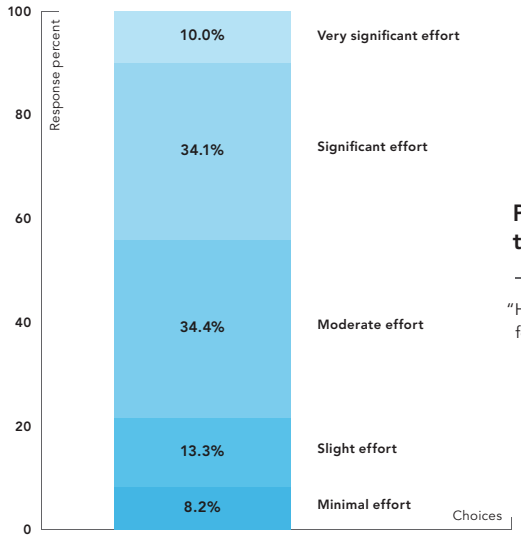

Perceived Efforts Required to Train LLMs for the Enterprise

"How significant is the effort you anticipate training LLMs for your business and deploy the model(s) effectively?"

Business and Technical Executives are being asked by their boards and increasingly by shareholders how they plan to leverage this new dawn of AI and the promise it provides to unlock language to solve problems. Around one-third of enterprises are already considering building enterprise-specific language models (37.1%). A significant majority of enterprises (78.5%) realize that the efforts required to effectively train a usable and accurate enterprise-specific language model is a major undertaking. Teams are already budgeting for LLM adoption and training projects, with 17.4% having available budget this year, another 17.7% planning to allocate budget and 39.5% discussing budgeting for next year.

This makes sense, as most of the public domain data used to train LLMs like ChatGPT is not enterprise-grade or domain-specific data. Even if a language model has been trained on different domains, it is not likely representative of what is used in most complex enterprise use cases, whether vertical domains like Financial Services, Insurance, Life Sciences and Healthcare, or highly specific use cases like contract review, medical claims, risk assessment, fraud detection and cyber policy review. Training effort will be required to have quality and consistent performance within highly specific domain use cases.

Generative AI and LLM Predictions for Natural Language Adoption

Given that language is the key currency of generative AI and LLMs, how will this impact how organizations are adopting natural language technologies?

"I think it will save considerable amount of time and serves as a great jumping off point for content creations but could be problematic when the content all sounds the same." **Business Stakeholder in Media & Publishing**

"Once the hype has calmed down (this is similar to previous technology roll outs) the development of practical use cases will prove difficult. At the same time 'private' large language models built in house will add additional biases to an organization causing a risk to the business form a spiral of hearing what you want to rather than what might actually be real." **Consultant in Insurance**

"Best case, as a partner for generating ideas. For example, biomedical research is an enterprise akin to 100 near-sighted examiners of a very large elephant, each examining their own small patch. It seems like AI might help here." **Business Stakeholder in Life Sciences**

"I see significant applications in customer care, in providing the first responses to customers, in classifying their requests. Other important applications in the analysis and classification of data contained in the company's digital archives." **C-Level Executive in Business Services**

"I believe that Gen AI and LLM will have a significant and potential disruptive impact if not managed carefully at the employee/student level. From the educational perspective it is now incumbent on faculty to raise the level of teaching done in the classroom to allow student to find the basic answers to questions through Gen AI and LLM." **Academic**

expert.ai

Summary and Considerations

Large Language Models are a hot topic in the field of artificial intelligence, and several models have gained prominence in recent years, such as GPT, LLAMA, Lambda, Anthropic and Ai21. These models have demonstrated impressive capabilities in natural language processing tasks, including text generation, translation and summarization. They also present significant adoption challenges that enterprises must address prior to widespread deployment and value realization. The LLM adoption path enterprises choose will likely vary by industry, use case and risk appetite. Regardless of the path, enterprises around the world are going to be increasingly leveraging generative AI and LLMs to improve productivity, augment current staff and drive competitive advantage.

Choose Appropriate Use Cases and Tools

Enterprises today understand that there's a lot of value hidden in all the unstructured data they handle on a daily basis. Many see natural language processing as one big technology and assume that, while vendors' models and tools might differ in product quality and price, ultimately they are largely the same. The truth is, NLP is not one thing; it's not one tool, but a toolbox. In the vast majority of situations, choosing which use cases to target should be a decision that is driven by the business to augment current team processes and eliminate low-value work. Teams often believe the myth that adopting LLMs in the enterprise requires large repositories of proprietary data for training. It's just not true. In fact, there are many ways to train language models to deal with very specific topics that may only have a few training resources available.

Include a Human Component

Having humans at only the beginning or only the end of an AI process is not enough to ensure accuracy, transparency or accountability. Enterprises need a human-centered approach, where data and inputs can be monitored and tuned by users throughout the process. Explainable-by-design and interpretable-by-design AI models offer humans control and the ability to transfer the process-specific knowledge that is needed to achieve the highest levels of accuracy. Also, keeping a human subject matter expert in the loop will significantly help with solution adoption and the fine tuning that is needed as data evolves.

Plan to Customize Your Models

Very few enterprises, if any, can invest in the creation, training and tuning of LLMs. Given that most enterprise use cases will require some degree of domain-specific training, there is a growing trend towards developing smaller LLMs. These smaller models are designed to be more efficient, faster and require less computational resources while maintaining high performance. Work with language experts to choose the best model and training approach.

Don't Forget About Governance Principles

In attempts to be early adopters with the latest and greatest technologies, enterprises have been known to put governance principles aside. When dealing with AI governance this is not the best approach. The risks are too high not to have an active and vigilant governance stance when it comes to AI concepts like data privacy, bias, truthfulness and responsible use. The potential of looming regulatory oversight suggests the need for executive-level participation in AI governance planning.

About expert.ai

Expert.ai (EXAI:IM) is a leading company in AI-based natural language software. Organizations in insurance, banking and finance, publishing, media and defense all rely on expert.ai to turn language into data, analyze and understand complex documents, accelerate intelligent process automation and improve decision making. Expert.ai's purpose-built natural language platform pairs simple and powerful tools with a proven hybrid AI approach that combines symbolic, machine learning, and Large Language Models LLMs to solve real-world problems and enhance business operations at speed and scale. With offices in Europe and North America, expert.ai serves global businesses such as AXA XL, Zurich Insurance Group, Generali, The Associated Press, Bloomberg, BNP Paribas, Rabobank, Gannett and EBSCO.
For more information, visit **https://www.expert.ai**.

Methodology

The survey results were collected from 300+ NLP practitioners, executives and academics from around the world. The interviews were conducted online by expert.ai in April 2023.