



2023

The large language model operations (LLMOps) market map



Where is all this data from?

**The CB Insights platform
has the underlying data
included in this report**

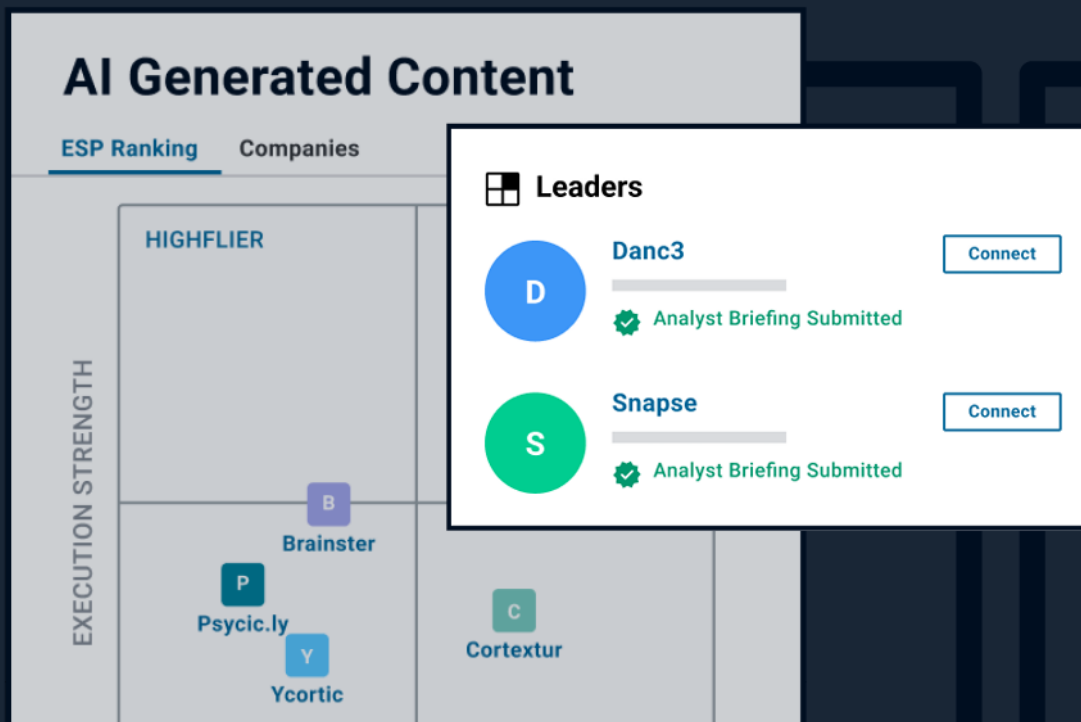
[Click here to sign up for free](#)



What is CB Insights?

CB Insights is a tech market intelligence platform that analyzes millions of data points on venture capital, startups, patents, partnerships and news mentions to help you see tomorrow's opportunities, today.

[Click here to learn more](#)



From prompt engineering to machine learning security, we break down the vendors helping enterprises deploy large language models.

The launch of OpenAI's ChatGPT sparked an unprecedented surge of corporate interest in large language models (LLMs).

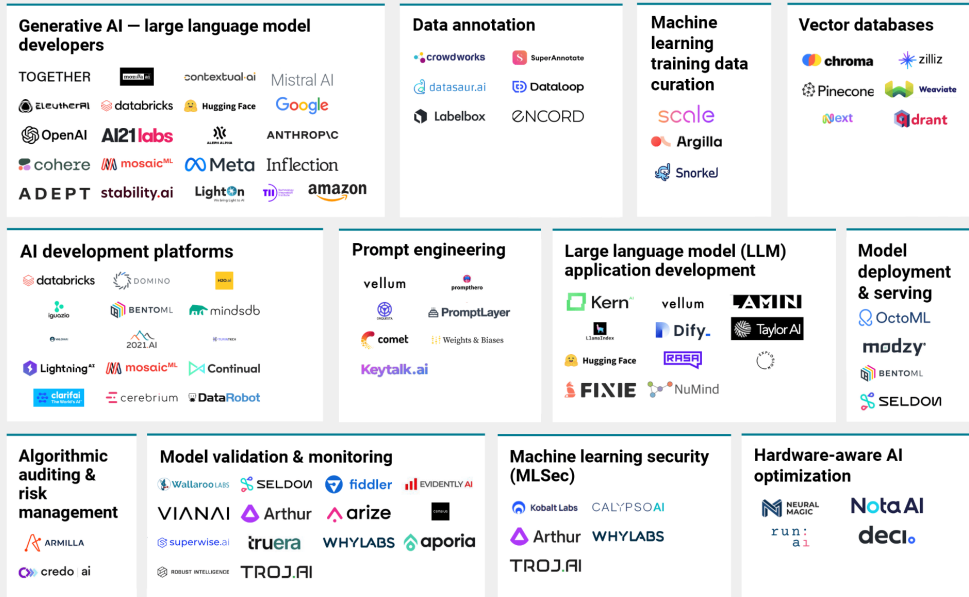
Big tech companies and research institutes are democratizing access to LLMs, enabling enterprises to more easily take advantage of deep learning capabilities. As enterprises adopt these models, they are focused on improving their data infrastructure, fine-tuning models for specific use cases, and monitoring LLMs for hallucination and bias.

As a result, tech vendors supporting large language model operations (LLMOps) – the end-to-end workflow that organizations employ to build, fine-tune, and deploy LLMs into production – are gaining traction.

In the market map below, we identify 90+ companies across 12 different categories helping enterprises bring LLM projects from start to finish.

Note: Our map includes public, private, and recently exited companies in LLMOps. This market map is not exhaustive of the space.


Large language model operations (LLMOps) market map



CBINSIGHTS

Market comparisons


Large language model operations (LLMOps) total equity funding



Source: CB Insights

CBINSIGHTS

Large language model operations (LLMOps) equity deal count



Source: CB Insights

CBINSIGHTS

Market descriptions

[Generative AI – large language model developers](#)

The generative AI – large language model developers market offers foundation models and APIs that enable enterprises to build natural language processing applications such as content creation, summarization, classification, chatbots, sentiment analysis, and more. Enterprises can fine-tune and customize these large-scale language models – pre-trained on vast amounts of text – for their specific use cases.

Equity funding 2023 YTD: \$15.3B|20 deals

Headcount 1-year change: -5%

Featured companies:

[Together](#) [Mozilla.ai](#) [Contextual AI](#) [Mistral AI](#) [EleutherAI](#) [Databricks](#) [Hugging Face](#) [Google](#) [OpenAI](#) [AI21 Labs](#) [Aleph Alpha](#) [Anthropic](#) [Cohere](#) [MosaicML](#) [Meta](#) [Inflection AI](#) [Adept](#) [Stability AI](#) [LightOn](#) [Amazon](#) [Technology Innovation Institute](#)

(A2A) payments ([Banked](#) and [Kevin](#)), and mobile wallets and remittances ([DANA](#) and [LemFi](#))

[Data annotation](#)

The data annotation market provides services for labeling large volumes of data in preparation for training AI and ML models. This market comprises both text and image & video annotation services. Most vendors employ human annotators to classify and label datasets, with some offering AI-powered automation tools to speed up the process.

Equity funding 2023 YTD: \$65M|3 deals

Headcount 1-year change: +20%

Featured companies:

[CrowdWorks](#) [SuperAnnotate](#) [Datasaur](#) [Dataloop](#) [Labelbox](#) [Encord](#)

Machine learning training data curation

The machine learning training data curation market offers solutions to support data quality control in the AI algorithm training process. These solutions help organizations complete key tasks, such as selecting the best subsets of data for training models, triaging datasets for bias, and identifying labeling errors. Ultimately, these solutions help minimize the downstream effects of poor-quality data on AI performance.

Equity funding 2023 YTD: \$2M|1 deal

Headcount 1-year change: +43%

Featured companies:

[Scale](#) [Argilla](#) [Snorkel AI](#)

Vector databases

The vector databases market focuses on providing databases optimized for high-dimensional, vector-based data. These databases are designed to efficiently store, manage, and query large volumes of vectors – i.e., mathematical representations of data points in multidimensional space.

Vector databases cater to a wide range of applications, including machine learning, natural language processing, recommendation systems, and similarity search.

Equity funding 2023 YTD: \$176M|5 deals

Headcount 1-year change: +68%

Featured companies:

[Chroma](#) [Zilliz](#) [Pinecone](#) [Weaviate](#) [NNext](#) [Qdrant](#)

AI development platforms

The AI development platforms market offers solutions that serve as one-stop shops for enterprises that want to develop and launch in-house AI projects. Vendors in this space enable organizations to manage all aspects of the AI lifecycle – from data preparation, training, and validation to model deployment and continuous monitoring – through a single platform in order to facilitate end-to-end model development. Some vendors offer “drag-and-drop” interfaces or “plug-and-play” solutions that enable teams without in-depth AI expertise to build AI projects.

Equity funding 2023 YTD: \$837M|11 deals

Headcount 1-year change: +21%

Featured companies:

[Databricks](#) [Domino](#) [H2O.ai](#) [Iguazio](#) [BentoML](#) [MindsDB](#) [Valohai](#) [2021.AI](#) [TurinTech](#)
[Lightning AI](#) [MosaicML](#) [Continual](#) [Clarifai](#) [Cerebrum](#) [DataRobot](#)

Prompt engineering

The prompt engineering market focuses on designing and generating specific prompts for machine learning models. By carefully crafting the input given to a language model, prompt engineering aims to guide the model toward generating accurate outputs that align with the user's intentions. This process involves understanding the strengths and weaknesses of the model and the task at hand. Ultimately, the solutions in this space aim to improve model performance and reliability.

Equity funding 2023 YTD: \$56M|3 deals

Headcount 1-year change: +18%

Featured companies:

[Vellum](#) [PromptHero](#) [Orquesta](#) [PromptLayer](#) [Comet](#) [Weights & Biases](#) [Keytalk AI](#)

Large language model (LLM) application development

The large language model (LLM) application development market includes tools for customizing and refining pre-trained language models for specific tasks and industries. Fine-tuning involves adjusting the weights of a model or training the model on task-specific data to make it more accurate and adaptable for particular applications. Companies in this market offer services and tools to fine-tune large language models like GPT-3 or open-source models.

Equity funding 2023 YTD: \$314M|11 deals

Headcount 1-year change: +36%

Featured companies:

[Kern AI](#) [Vellum](#) [Lamini](#) [LlamaIndex](#) [Dify](#) [Taylor AI](#) [Hugging Face](#) [Rasa](#) [Explosion](#) [Fixie.ai](#)
[NuMind](#)

Model deployment & serving

The model deployment & serving market bridges the gap between data science and DevOps teams by taking trained machine learning models and putting them into production. Vendors offer tools for machine learning deployment on Kubernetes as well as serverless technology that can be used to deploy AI in cloud and on-prem environments. Most deployment vendors provide continuous model monitoring and governance tools.

Equity funding 2023 YTD: \$29M | 2 deals

Headcount 1-year change: -7%

Featured companies:

[OctoML](#) [Modzy](#) [Seldon](#) [BentoML](#)

Algorithmic auditing & risk management

The algorithmic auditing & risk management market provides solutions for evaluating and mitigating risks associated with algorithmic decision-making. These tools enable organizations to ensure algorithmic fairness, transparency, and regulatory compliance. Vendors in this space take a multifaceted approach to derisking AI, which includes data auditing, model validation, metadata tracking, and post-production monitoring.

Equity funding 2023 YTD: \$1M|2 deals

Headcount 1-year change: +22%

Featured companies:

[Armillia AI](#) [Credo AI](#)

Model validation & monitoring

The model validation & monitoring market provides solutions that continuously monitor the performance of AI models and provide real-time visibility into model behavior. These solutions track outliers in predictions, potentially biased outcomes, and suspected adversarial attacks. Demand for these solutions is driven by the fact that AI model performance can degrade over time if it continuously encounters real-world data that varies significantly from its training data.

Equity funding 2023 YTD: \$20M|3 deals

Headcount 1-year change: +13%

Featured companies:

[Wallaroo](#) [Seldon](#) [Fiddler AI](#) [Vianai Systems](#) [Arthur](#) [Arize](#) [Superwise](#) [TruEra](#) [WhyLabs](#)
[Robust Intelligence](#) [TrojAI](#) [Aporia](#) [Evidently AI](#) [Censius](#)

Machine learning security (MLSec)

The machine learning security (MLSec) market provides solutions designed to protect machine learning models and algorithms from

adversarial attacks, data poisoning, model evasion, backdoor injections, and other cyber attacks. Vendors offer a range of products, including intrusion detection systems, adversarial defense systems, secure machine learning frameworks, and anomaly detection tools.

Equity funding 2023 YTD: \$109M | 4 deals

Headcount 1-year change: +35%

Featured companies:

[Kobalt Labs](#) [Calypso AI](#) [Arthur](#) [WhyLabs](#) [TrojAI](#)

Hardware-aware AI optimization

The hardware-aware AI optimization market provides software solutions that optimize AI algorithms and models to run efficiently on available hardware, such as GPUs and CPUs. These solutions also allow enterprises to compress neural networks to run on edge devices or on-prem servers. With optimization tools, businesses can speed up AI deployments, reduce prediction latency, and improve model performance.

Equity funding 2023 YTD: No deals

Headcount 1-year change: +22%

Featured companies:

[Neural Magic](#) [Run:AI](#) [Nota AI](#) [Deci](#)

If you aren't already a client, [sign up for a free trial](#) to learn more about our platform.