


 EMERGING TECH RESEARCH

Artificial Intelligence & Machine Learning Report

VC trends and emerging opportunities

Q1
2023





Contents

Vertical overview	3
Q1 2023 timeline	4
AI & ML landscape	5
Generative AI market map	6
VC activity	7
Emerging opportunities	15
Generative AI applications	16
Large language model operations	19
Generative AI semiconductors	22
Select company highlights	24
Databricks	25
DataRobot	27

Institutional Research Group

Analysis



Brendan Burke Senior Analyst, Emerging Technology
brendan.burke@pitchbook.com
pbinstitutionalresearch@pitchbook.com

Data

Matthew Nacionales Data Analyst

Publishing

Report designed by **Chloe Ladwig** and **Julia Midkiff**

Published on June 29, 2023



For previous updates as well as our complete AI & ML research, please see the designated [analyst workspace](#) on the PitchBook Platform.




Vertical overview

The AI industry is at an inflection point between pre- and post-LLM (large language model) eras. The recent wave of LLMs has subsumed end-user attention, forcing incumbents to expand their product portfolios to continue revenue growth. Leading tech giants such as [Microsoft](#), [Amazon](#), [Google](#), [Meta](#), Nvidia, [Intel](#), [Qualcomm](#), [HPE](#), and [Salesforce](#) are making a series of prodigious corporate investments in both internal and external technologies, benefiting research-focused startups and academic researchers. Each of these companies has participated in multiple VC megadeals since the start of 2022, and their participation is a key component of the largest deals in the space. Additionally, [Databricks](#) and [Snowflake](#) have recently acquired AI startups to accelerate their competition. Even so, the last mile of application building remains a “Wild West” that can support a range of services companies and infrastructure startups.

Open-source growth offers a path to market for startups. AI projects are achieving record open-source community growth in terms of [GitHub](#) star count, indicating the number of developers paying attention to the projects. [Stability AI](#)’s Stable Diffusion model grew rapidly in open-source interest last year, yet recent language model releases have outpaced its user growth, reaching 15,000 stars almost immediately upon launch. These include forks of [Meta](#)’s leaked LLM LLaMa, made available via Stanford with its Vicuna model. For comparison, leading Apache projects only garner 20,000-25,000 stars yet have produced large enough developer communities to support high-value companies such as [Databricks](#) and [Confluent](#). Generative AI may not yet be producing significant revenue totals for startups, yet developer interest is the clearest path to commercialization for emerging architectures and supporting tools.


GitHub star count for selected open-source LLM repositories



Source: [Star History](#) • Geography: Global • *As of June 8, 2023



Q1 2023 timeline



Q1 VC activity

1,497

total deals

\$22.7B

total VC raised

43.0%

deal value growth QoQ

2023 YTD summary

1,497

total deals

\$22.7B

total VC raised

-20.2%

deal value growth YoY



AI & ML landscape

- 1 Horizontal platforms
- 2 Vertical applications
- 3 Semiconductors
- 4 Autonomous machines





AI & ML VC ecosystem market map

Click to view the interactive market map on the PitchBook Platform.

Market map is a representative overview of venture-backed or growth-stage providers in each segment. Companies listed have received venture capital or other notable private investments.

Horizontal platforms

AI automation platforms



AI core



Computer vision



Natural language technology



Vertical applications

Consumer




Financial services



Healthcare



Industrial



Information technology



Transportation



AI & ML semiconductors

Edge AI software



Intelligent sensors & devices




Processor design



Autonomous machines

Autonomous vehicles



Intelligent robotics






VC activity

VC activity ticked up in Q1 based on [OpenAI](#)'s unconventional corporate investment from [Microsoft](#). We tracked \$22.7 billion in investment during the quarter, including \$12.7 billion from generative AI investments including [OpenAI](#). If [OpenAI](#)'s deal is excluded, then VC deal activity continued to decline in terms of deal value and count. Vertical applications companies only raised \$6.0 billion in Q1 after raising \$50.6 billion in 2022, demonstrating weakness in conventional domain-specific software. Investors lack interest in software solutions unless they drive the generative field forward. Generative AI software continues to earn outstanding valuation multiples, with research startups striking deals based on the level of computing power and talent they will need to make breakthroughs. VC investors are adapting to these requests yet not supporting the capital demands of more mature companies that are unlikely to make research breakthroughs.


We are beginning to see the beginning of a generative AI acquisition trend. [Microsoft](#)'s investment in [OpenAI](#) is structured similarly to an acquisition, with a reported profit share and exclusive distribution rights, demonstrating the value of third-party technology to tech giants. Both [Snowflake](#) and [Databricks](#) have recently acquired generative AI startups in [Neeva](#) and [Rubicon](#), respectively. [Neeva](#) is a neural search startup with significant latency improvements over conventional AI searches, while [Rubicon](#) focuses on data storage. Apart from these trends, acquisition values have been minimal, only reaching \$3.3 billion in disclosed deal value in Q1 after reaching \$48.7 billion in disclosed exit value in 2022. Tech giants were mostly dormant, and [Iguazio](#)'s \$50.0 million acquisition by [McKinsey](#) fell below the company's prior private valuation.

AI & ML VC deal activity



Source: PitchBook • Geography: Global • *As of March 31, 2023

AI & ML VC exit value (\$B) by type




Source: PitchBook • Geography: Global • *As of March 31, 2023




VC ACTIVITY

Median AI & ML VC deal value (\$M) by stage




Source: PitchBook • Geography: Global • *As of March 31, 2023

Median AI & ML VC pre-money valuation (\$M) by stage




Source: PitchBook • Geography: Global • *As of March 31, 2023

AI & ML VC deal value (\$M) by vertical application



Source: PitchBook • Geography: Global • *As of March 31, 2023

AI & ML VC exit count by type



Source: PitchBook • Geography: Global • *As of March 31, 2023



VC ACTIVITY

Key AI & ML early-stage VC deals by deal value*

Company	Close date	Subsegment	Deal value (\$M)	Pre-money valuation (\$M)	Lead investor(s)
Abound	March 6, 2023	AI in financial services	\$602.0	N/A	Hambro Perks, K3 Ventures, GSR Ventures
Adept	February 23, 2023	AI core	\$350.0	\$650.0	General Catalyst, Spark Capital
ONE	February 1, 2023	Autonomous vehicles	\$300.0	\$900.0	Fifth Wall, Franklin Templeton
Cohere	March 28, 2023	Natural language technology	\$250.0	\$1,750.0	Inovia Capital
MiniMax AI	January 9, 2023	Consumer AI	\$250.0	\$950.0	MiHoYo, Future Capital Discovery Fund, Tencent
Character.ai	March 7, 2023	Consumer AI	\$150.0	\$700.0	Andreessen Horowitz
Deepway	March 20, 2023	Autonomous vehicles	\$111.4	N/A	SB China Venture Capital, Weiqiao Pioneering Group
Pasqal	January 17, 2023	AI in healthcare	\$103.2	\$402.4	Temasek Holdings
Artera	March 21, 2023	AI in healthcare	\$90.0	N/A	Johnson & Johnson Innovation - JJDC, Coatue Management, Koch Disruptive Technologies
Vannevar Labs	January 18, 2023	Natural language technology	\$75.0	\$500.0	Felicis

Source: PitchBook • Geography: Global • *As of March 31, 2023



VC ACTIVITY

Key AI & ML late-stage VC deals by deal value*

Company	Close date	Subsegment	Stage	Deal value (\$M)	Lead investor(s)	Valuation step-up
OpenAI	January 23, 2023	AI core	N/A	\$10,000.0	Microsoft	N/A
Monogram Health	January 9, 2023	AI automation platforms	Series C	\$375.0	CVS Health, Cigna, Frist Cressey Ventures, Humana Memorial Hermann Foundation, Norwest Venture Partners, TPG	N/A
Asimov	January 5, 2023	AI in healthcare	Series B	\$175.0	CPP Investments	9.5x
Amelia	March 7, 2023	Natural language technology	N/A	\$175.0	BuildGroup, Monroe Capital	N/A
SirionLabs	January 4, 2023	AI automation platforms	Series D	\$110.0	Partners Group	N/A
Trusting Social	February 10, 2023	AI in financial services	Series D	\$105.0	Masan Group	N/A
Humane	March 8, 2023	Intelligent sensors & devices	Series C	\$100.0	Kindred Ventures	1.1x
Zhida Network	March 28, 2023	Transportation	Series C2	\$84.0	Yixin Group, Chengdu Technology Transfer Venture, Chengdu Tongchuang Zhixing, Zhuhai Haidu Science and Technology Investment, Tencent Investment	N/A
BRETON Technology	January 5, 2023	Autonomous vehicles	Series C1	\$83.2	Inner Mongolia Eerduosi Resources Company, Hunan Chasing Fund	N/A
Cardinal Operations	March 31, 2023	AI automation platforms	Series C2	\$72.5	National Manufacturing Transformation and Upgrade Fund	N/A

Source: PitchBook • Geography: Global • *As of March 31, 2023



VC ACTIVITY

Key AI & ML VC exits by exit value (\$M)*

Company	Close date	Segment	Category	Exit value (\$M)	Acquirers(s)/Index	MOIC
Hesai Technology	February 9, 2023	Vertical applications	Transportation	\$2,201.8	NASDAQ	4.0x
AnyMind Group	March 29, 2023	Vertical applications	Consumer AI	\$419.1	Tokyo Stock Exchange	4.2x
GEO Semiconductor	March 3, 2023	AI & ML semiconductors	Processor design	\$270.0	indie Semiconductor	3.5x
Gaivota	February 24, 2023	Vertical applications	Industrial AI	\$235.0	Seedz Cayman	N/A
Iguazio	January 23, 2023	Horizontal platforms	AI core	\$50.0	McKinsey & Company	0.6x
UntieNots	January 12, 2023	Vertical applications	AI in IT	\$37.9	Eagle Eye Solutions Group	10.5x
Cipher	March 31, 2023	Vertical applications	AI in IT	\$36.8	RELX	2.0x
DW	February 15, 2023	Horizontal platforms	AI automation platforms	\$36.7	Impress	47.9x
Intechnica	January 17, 2023	Vertical applications	AI in IT	\$17.6	Crosslake Technologies, Falfurrias Capital Partners	N/A
Hongyaa	January 16, 2023	Vertical applications	Consumer AI	\$11.6	Seentao Technology	8.2x

Source: PitchBook • Geography: Global • *As of March 31, 2023



VC ACTIVITY

Top strategic acquirers of AI & ML companies since 2017*

Acquirer	Deal count	Investor type
Apple	21	Corporation
Accenture	19	Corporation
Microsoft	14	Corporation
Meta Platforms	12	Corporation
Cisco Systems	11	Corporation
ServiceNow	10	Corporation
DataRobot	9	VC-backed company
Intel	8	Corporation
IBM	8	Corporation
Oracle	7	Corporation

Source: PitchBook • Geography: Global • *As of March 31, 2023



VC ACTIVITY

Top VC investors in AI & ML companies since 2017*

Investor	Deal count	Angel and seed	Early-stage VC	Late-stage VC	Venture growth	Investor type
Alumni Ventures	315	129	98	75	13	VC
Sequoia Capital China	259	26	160	58	15	VC
SOSV	229	105	58	58	8	VC
Sequoia Capital	210	51	78	55	26	VC
500 Global	209	97	69	38	5	VC
Right Side Capital Management	200	94	92	14	0	VC
Tiger Global Management	192	4	57	89	42	VC
Khosla Ventures	192	57	64	55	16	VC
Accel	191	36	76	64	15	VC
Andreessen Horowitz	186	31	76	56	23	VC

Source: PitchBook • Geography: Global • *As of March 31, 2023



VC ACTIVITY

Top VC-backed AI & ML companies by total VC raised to date*

Company	VC (\$M) raised to date	Segment	Category	HQ location
OpenAI	\$11,010.1	Horizontal platforms	AI core	San Francisco, US
Waymo	\$5,500.0	Autonomous machines	Autonomous vehicles	Mountain View, US
Databricks	\$3,497.4	Horizontal platforms	AI automation platforms	San Francisco, US
JD Digits	\$3,141.3	Vertical applications	AI in financial services	Beijing, China
Weltmeister	\$2,749.6	Autonomous machines	Autonomous vehicles	Shanghai, China
Anduril	\$2,315.1	Autonomous machines	Intelligent robotics	Costa Mesa, US
Nuro	\$2,132.0	Autonomous machines	Intelligent robotics	Mountain View, US
Horizon Robotics	\$1,900.0	AI & ML semiconductors	Chips	Beijing, China
UBTECH	\$1,861.0	Autonomous machines	Intelligent robotics	Shenzhen, China

Source: PitchBook • Geography: Global • *As of March 31, 2023



Emerging opportunities

Generative AI applications

Developer adoption of code-completion tools makes the space likely to stand out in commercial results.

Large language model operations

LLM implementation requires a custom stack of software solutions to take foundation models to production for ongoing inference on user data.

Generative AI semiconductors

AI chip startups exhibit software innovation that may yield sustainable market positions.



Generative AI applications

Overview


The generative AI landscape has continued to evolve rapidly since our [Vertical Snapshot](#) published in Q1. Increasingly, multimodal foundation models offer possibilities for dialogue with AI models about both text and images, while progress is made on video generation. We have tracked \$12.7 billion invested in generative AI through Q1, including [OpenAI](#)’s outlier investment. This total includes companies developing and commercializing the most recent developments in transformer and diffusion architecture for language and image generation. It does not include end users of the technology or all those that have added a product feature including generative models.

We have added critical use cases to our taxonomy including fintech, edtech, and code testing. Fintech was a focus of the recent YCombinator class, particularly workflow solutions for finance professionals. Edtech stands out as an obvious use case of LLMs given their usage in schools to date. Developer tools are proliferating beyond code completion as models add the ability to make context-driven responses for documentation and testing. Additionally, we have broken out the infrastructure layer (see section on LLM operations on page 19) given the rapid stratification of layers of the LLM stack.

Market direction

Breakthroughs can most directly have an impact in code generation. Code generation leads all use cases in VC funding so far in 2023, raising \$409.7 million, on pace to surpass 2022’s total. Bolstering this result is the high adoption from developers of coding assistants and the rapidly improving performance of foundation models on coding tasks. Some surveys suggest that 90% of developers have used ChatGPT. Looking at commercial products, 70% of developers have

Generative AI vertical application market map



Note: Excluding code, biotech, visual, audio, and language



GENERATIVE AI APPLICATIONS


used [GitHub](#)’s Copilot, while leading VC-backed solutions have gained around 20% penetration. Underlying Copilot’s success, GPT-4 made breakthroughs in code quality based on the inclusion of [Microsoft’s GitHub](#) data. Developers have encountered trouble exposing sensitive code to cloud-based products. As a result, open-source models specifically for code generation are emerging to supplement [OpenAI](#), including WizardLM and [Meta](#)’s leaked LLaMa model.

Large enterprises are struggling to adopt generative applications at scale given privacy concerns and accuracy needs. [Stability AI](#) has faced press reports that the company is falling short of revenue targets and fundraising goals, even after disclosing pilot projects with automakers and movie studios for custom model development.¹ While these struggles are natural for a new enterprise, investors may have expected faster commercialization when granting unicorn valuations at an early stage, creating challenges around follow-on funding. [Character AI](#) demonstrates that viral adoption is continuing for AI products, having nearly tripled its website traffic from January to April, amounting to 173.1 million visitors, according to our data from Similarweb.² Staggering usage of free tools is not rapidly translating into paid pilot projects for startups yet continues to achieve records for public interest in AI.

Trending startups

Pre-LLM startups are benefiting from new techniques to grow their user bases and valuations. In April 2023, coding assistance [Replit](#) reached a \$1.2 billion post-money valuation in its Series B1 led by Andreessen Horowitz. The company has benefited from high adoption in developing countries,

1: “Stability AI Is on Shaky Ground as It Burns Through Cash and Looks at a Management Overhaul,” Semafor, Reed Albergotti, April 7, 2023.
2: “Character.ai Traffic and Engagement Analysis,” Similarweb, May 2023.



Source: IDC • Geography: Global • *As of March 31, 2023




GENERATIVE AI APPLICATIONS

showing the growth markets that can be expanded with new technology. Generative AI slide presentation software startup [Tome](#) raised a \$43.0 million Series B with a 1.5x valuation step-up while disclosing the fastest growth of any productivity tool to 1 million users. [Tome](#) embraces the command line interface as a creative starting point for generative AI, as users can type in the type of the presentation they would like to create, then receive a menu of common customizations that are often manually edited in PowerPoint. The founders Keith Peiris and Henri Liriani previously worked on Instagram and Facebook Messenger, enabling them to stand out in delivery of AI products via user interfaces. The company relies on external AI models including Stable Diffusion XL and GPT-4, enabling multimodal outputs within the same project.

YCombinator startups brought generative AI to a range of applications yet focused on areas with positive customer response, including sales and finance. Of the 45 companies graduating the Winter 2023 batch, nine focused on sales, serving as by far the most popular use case. We observed outstanding commercial traction for startups in this space, including [Fabius](#), [Sameday](#), and [Sail](#). Other novel use cases included dubbing for YouTube content creators ([EzDubs](#)), Python-integrated spreadsheets ([Neptyne](#)), generative studio-quality video ([Linum](#)), and industry research reports ([AiFlow](#)). The Spring 2023 cohort is still taking shape yet seems likely to maintain a focus on generative AI.

Generative AI VC deal activity



Source: PitchBook Emerging Tech Research • Geography: Global • *As of March 31, 2023



Large language model operations

Overview

LLM implementation requires a custom stack of software solutions to take foundation models to production for ongoing inference on user data. While the power of foundation models materialized in Q1, the challenges of deploying them in customer environments drove the field in subsequent months. Developers experimenting with foundation models leaned on new packages of software scripts to automate critical steps of LLM application development, encouraging the adoption of “picks and shovels” of AI applications.

Customized software solutions earn high developer usage across LLM architecture, orchestration, and deployment. In model architecture, startups not only train novel models but also optimize model outputs. Reinforcement learning from human feedback (RLHF) has emerged as a critical service that can be offered as a service. [OpenAI](#) worked with data preparation startup [Sama](#) for data labeling of its model outputs, which proved critical for the reliability of its ChatGPT product. Data preparation unicorn [Scale AI](#) works with all leading LLM research labs on RLHF. In orchestration, startups connect user databases to model predictions via prompt engineering and retraining of deployment-ready model versions. In deployment, companies compress AI models and collect metrics on their performance. While serving models to cloud hardware has several valid approaches, performance monitoring is an open research field with new techniques emerging to assess accuracy and reliability.

Foundation LLM operations critical capabilities



Source: PitchBook • Geography: Global




LARGE LANGUAGE MODEL OPERATIONS

Market direction

Large developer communities are forming around open-source projects that can support high-revenue businesses over time. To connect user databases to language models, developers have utilized customized software from open-source libraries such as [LangChain](#). [LangChain](#) features prebuilt scripts for building chatbots and search tools with leading foundation models. The company faces controversy from developers for lack of generalization to new use cases and unfamiliar data types yet has attracted a strong developer community, encouraging contributions and regular releases of new integrations and functions including agentic AI. A wide range of fine-tuning solutions have emerged to solve the “last mile.” We are tracking 30 companies doing some form of fine-tuning, yet only 14 have raised VC capital to date. Nine such companies have recently graduated from YCombinator: [Anarchy](#), [BerriAI](#), Chima, Defog, [Gloo](#), [GradientJ](#), [Magicflow](#), [Psychic](#), and [Pyq](#).

Going forward, agents, cybersecurity, and monitoring tools will be needed to make foundation models into time savers for enterprise workers and create virtual assistants. Agents refer to software bots that can receive instructions from LLMs and use software connectors to act in different contexts to implement the outcome. Examples of these agents include sending text messages to services to embed or code samples to third-party services to run. These functions are driven by open-source projects such as AgentGPT, BabyAGI, and AutoGPT, each of which is attracting massive developer numbers, as measured by [GitHub](#) stars, that exceed the totals of iconic open-source projects such as Apache Spark and Kafka, commercialized by [Databricks](#) and [Confluent](#), respectively. Pre-LLM monitoring solutions can still be used to track the outputs of LLMs with modified measurement techniques, while novel cybersecurity solutions are emerging to detect novel attacks on new models. Popular open-source projects will be soon followed by companies commercializing them and raising outsized rounds of funding.

Foundation LLM operations market map





LARGE LANGUAGE MODEL OPERATIONS


Trending startups

Vector database startups are converting open-source adoption into significant funding rounds to build a moat around their LLM services. In April, [Chroma](#), [Qdrant](#), [Weaviate](#), and [Pinecone](#) raised rounds at high valuations. [Pinecone](#) achieved a 3.9x valuation step-up. [Pinecone](#) contributed to founding this niche in 2021 and has grown its customer logo count by 8x to around 1,600 in the past three months, according to Andreessen Horowitz.³ While the need for a litany of vector databases is controversial, we believe these startups can differentiate themselves based on the speed of AI inference using their databases, with [Qdrant](#) standing out for coding its databases in hardware-friendly coding language Rust and achieving high-end results in inference speed.

Early-stage startups can break new ground in fine-tuning and AI agents. Fine-tuning [LangChain](#) has built mindshare in the AI developer community with its robust library of LLM scripts and emerging library of agents built on top of models, leading to a \$10.0 million seed round with a \$45.0 million post-money valuation in Q1. Recently, [Blueprint AI](#), co-founded by former MLOps startup [Algorithmia](#) founder Diego Oppenheimer, along with multiple leaders of Madrona Venture Labs, emerged out of stealth to build a library of AI agent bots with pre-seed funding from leading AI investors. Further, Greylock invested \$8.5 million in pre-seed startup [LlamaIndex](#), founded by a former ML engineer from cybersecurity startup [Robust Intelligence](#). [LlamaIndex](#) offers alternatives to vector databases for language data storage, including graphs and indexes. Top AI talent is gravitating to opportunities for new ventures, leading to a pipeline of pre-seed opportunities.

3: "Investing in Pinecone," Andreessen Horowitz, Peter Levine, et al., April 27, 2023.

Foundation LLM operations early-stage VC deal activity



Source: PitchBook • Geography: Global • *As of March 31, 2023




Generative AI semiconductors

Overview

Semiconductors stand to benefit as a primary outlet for generative AI spending. With the breakthrough in model quality from ChatGPT, cloud buyers and AI research labs accelerated their capital expenditure plans, spurring Nvidia to increase its revenue guidance above expectations by 55.5%, or \$3.9 billion, for Q2 2023. Even so, Nvidia does not own the entire AI value chain, and VC investors have already begun shifting their focus to categories in which startups can carve out market share. Hyperscalers' custom chips and startups can outperform the chip giant on specific inference tasks that will increase in importance as LLMs are rolled out from cloud data centers to customer environments.

We estimate the market for AI semiconductors to exhibit similar growth, surpassing Nvidia's growth expectations for this next quarter. The market for AI-specific data-center semiconductors is on pace to reach \$8.5 billion in 2023, and we forecast the market will reach \$23.6 billion by 2026. Graphical processing units such as Nvidia's make up the large majority of 2023's total but will contribute less than half of spending in 2026, ceding ground to smaller chips and application-specific processors. We believe application-specific AI chips can earn \$3.0 billion by 2026, which can support large businesses apart from Nvidia. The current small market size for these chips creates limits to scale for startup challengers.

AI data center semiconductor market size estimate (\$B)*



Source: PitchBook Emerging Tech Research • Geography: Global • *As of March 31, 2023



GENERATIVE AI SEMICONDUCTORS

Market direction

Chip startups can become verticalized computing platforms for specific model types. Startups have struggled to stay ahead of Nvidia's CUDA software when compiling AI models to hardware given the range of different frameworks. Generative AI offers an opportunity to innovate around a single family of transformer models for a range of tasks and optimize hardware. Further benefiting startups is developer consensus around PyTorch as the primary framework for training new AI models. Already, chip startups competed favorably with Nvidia for BERT large language models. [SambaNova Systems](#) disclosed significant speedups over Nvidia hardware for BERT training and inference, with the company's servers particularly shining in inference latency reduction with a 5.8x latency reduction.⁴ This advantage compounds if startups can continue to specialize in foundation models.

AI chip startups exhibit software innovation that may yield sustainable market positions. Based on a custom cluster of its wafer-scale processor, chip unicorn Cerebras trained open-source models based on a custom cluster of its semiconductors. The models are available on [Hugging Face](#). The open-source models have not achieved independently outstanding performance nor adoption yet may find use cases in the medium-term, offering a revenue stream for the company's processors. Cerebras has also partnered with generative AI unicorn [Jasper](#) to train custom models. [SambaNova Systems](#) has debuted a generative AI platform and partnered with the top AI service provider by market share, [Accenture](#). [Accenture](#) promotes the company's benefits in privacy and security. These startups have the potential to capture market share from Nvidia in customer-sited AI implementation.

4: "Breakthrough Efficiency in NLP Model Deployment," [SambaNova Systems](#), September 20, 2020.

Trending startups

Interconnect requirements for generative AI processors encourage the adoption of photonic processors, accelerating the maturation of this emerging field. In May, leading silicon photonics startup [Lightmatter](#) raised a \$153.3 million VC round. [Lightmatter](#) has optimized its hardware to conduct inference for leading neural network types including GPT-4. Also in May, [Ayar Labs](#) extended its Series C with a \$25.0 million investment from investors including Nvidia. Nvidia recognizes the need for high-speed interconnect between the assemblies of chips used for AI training. This investment also serves to future-proof Nvidia's business against photonic chips capturing market share from silicon in the long run.

At the seed stage, in April 2023, [Etched.ai](#) raised a \$5.4 million seed round, reaching a \$34.0 million post-money valuation—the highest we have tracked for an AI chip startup outside of China. The company is focusing on specialized chips for AI language model inference. The company uses high-bandwidth memory and high-speed interconnects to accelerate processing in LLM's vast latent space, claiming to outperform Nvidia's A100s by 140x on GPT-3 inference. Additionally, custom ASIC startup [Encharge AI](#) is expanding from computer vision use cases to LLM inference based on innovative in-memory computing, benefiting from interest from telecommunications companies in co-located data center AI processing. The company raised a \$21.7 million Series A in Q4 2022 based on the expectations for the company to deploy AI inference closer to the network edge yet still in data centers.



Select company highlights



SELECT COMPANY HIGHLIGHTS: DATABRICKS



Founded
2013

5,000 employees
in 25 offices globally

Total raised:
\$3.5B

February 2020:
Leader in Gartner Magic Quadrant
for Data Science and Machine
Learning Platforms

Last financing valuation:
\$38.0B

Last financing:
Raised \$1.6B in a Series H

Lead investors:
Counterpoint Global,
Franklin Templeton
Investments, Andreessen
Horowitz, NEA

Overview

Founded in 2013, [Databricks](#) offers a data science platform that includes AI-as-a-service (AlaaS) functionality with a suite of data science tools for data engineering, data warehousing, and ML algorithms. The company grew out of the open-source Apache Spark data science community and created an extensible Unified Data Analytics Platform that can ingest data from enterprise silos and prepare it for cluster-based computing. Once an effective open-source product was in place, the company moved to a closed-source model and rapidly increased revenue beginning in 2016.

[Databricks](#) now offers a product suite on top of Apache Spark that includes notebooks for ML model collaboration; a data lake, which integrates unstructured data in a central database for analytics; and data security. The company has focused on the generative AI opportunity, training a custom LLM called Dolly that can run more efficiently than GPT-4 and achieve competitive results on question-answering tasks. The model is open-source and rapidly grew to 10,000 [GitHub](#) stars. The company is adding to its infrastructure capabilities as well by acquiring AI data storage startup [Rubicon](#), hardware acceleration startup [MosaicML](#) for \$1.3 billion, and AI data governance startup [Okera](#). We believe the company is surpassing its peer [Snowflake](#) in AI capabilities.

Leadership

[Databricks](#)’ management has high technical expertise capable of generating product-led growth and has increased its public company experience. The company was founded by seven early contributors to Apache Spark, and all the cofounders remain at the company. CEO and co-founder Ali Ghodsi earned a Ph.D. in distributed computing and was VP of engineering before being promoted to CEO at the start of the company’s growth phase in 2016. CFO Dave Conte previously took Splunk public in the same role, which suggests [Databricks](#)’ accounts will be ready for scrutiny by public markets. The board features cofounders and VC investors, including Andreessen Horowitz and cofounder Ben Horowitz. In Q2 2022, [Databricks](#) hired the general counsel from DocuSign, which will provide public company experience as [Databricks](#) pursues a public listing.

Competitors

[Databricks](#) stands out in the AI & ML market by supporting machine learning operations (MLOps) atop its data lake, functionality that distinguishes it from legacy vendors such as [SAP](#), [Microsoft](#), and [AWS](#), while also being more AI-oriented than next-generation data warehouse vendor [Snowflake](#). Because of its ability to support batch processing for ML analytics and other high-value



SELECT COMPANY HIGHLIGHTS: DATABRICKS

workloads, [Databricks](#) is capturing market share from incumbents such as [Alteryx](#), [IBM](#), [Oracle](#), and [SAS](#). [Databricks](#) runs on [Microsoft](#) Azure and [AWS](#), although those hyperscalers would likely prefer to natively support their ML workflows and may improve their offerings in the future. The company gains market credit for scaling ML model deployment and management across a variety of compute clusters—a critical problem for data scientists new to ML.

Outlook

[Databricks](#) has achieved the necessary scale and growth to realize high performance in public markets. The company reached \$1.0 billion in revenue in 2022 with 60% revenue growth.

The company has cut its internal valuation by 10%, according to Ghodsi, suggesting that its valuation is staying relatively resilient as well.⁵ The company will need to defend itself against ML-optimized challengers that can carry out more efficient streaming data analysis and in-memory analytics within its database. Furthermore, it will benefit from analytics built on its database becoming essential to high-value industries, including IT and industrial. The company remains innovative and active in future-proofing its business, leading us to believe that, in the long run, [Databricks](#) will be valued more highly than relational database incumbents based on the growth opportunities of data lake architecture for streaming data.

5: “What Is Data Management Company Databricks Really Worth?” Fortune, Jessica Matthews, March 8, 2023.

Financing history

Series A	Series B	Series C	Series D	Series E	Series F
September 24, 2013	June 30, 2014	December 15, 2016	August 22, 2017	February 5, 2019	October 22, 2019
Total raised: \$14.0M	Total raised: \$33.0M	Total raised: \$60.0M	Total raised: \$140.0M	Total raised: \$250.0M	Total raised: \$400.0M
Pre-money valuation: \$34.0M	Pre-money valuation: \$214.0M	Pre-money valuation: \$453.0M	Pre-money valuation: \$800.0M	Pre-money valuation: \$2.5B	Pre-money valuation: \$5.8B
Investor(s): Andreessen Horowitz	Investor(s): NEA	Investor(s): NEA	Investor(s): Andreessen Horowitz	Investor(s): Andreessen Horowitz	Investor(s): Andreessen Horowitz



SELECT COMPANY HIGHLIGHTS: DATAROBOT



Founded 2012	Total raised: \$1.1B	Last financing valuation: \$6.3B	Lead investors: Tiger Global Management, Altimeter Capital Management, Sapphire Ventures, Meritech Capital Partners, NEA, Atlas Venture
1,265 employees in 8 offices globally	March 2021: Visionary in Gartner Magic Quadrant for Data Science and Machine Learning Platforms	Last financing: Raised \$300.0M in a Series G	

Overview

[DataRobot](#)’s Enterprise AI platform offers end-to-end automated machine learning (autoML) tooling, including data preparation, model development, security, and monitoring. The platform is designed for data scientists, software developers, operations teams, business line users, and executives. The platform can train and deploy models across cloud, on-premises, and hybrid production environments. Recent acquisitions in ML deployment ([Algorithmia](#)), autoML (Zeff), model management (ParallelIM), data visualization (Zepl), and data preparation (Paxata and Zeff) have been integrated to create a data science platform that can compete with hyperscalers.

According to management, the company was on pace for \$210.0 million in revenue in 2022.⁶ We believe that its revenue growth this year lags high-growth expectations, as we estimate the company is growing between 20% and 25%. AutoML has yet to become a large market relative to other enterprise software categories, which potentially constrains [DataRobot](#)’s achievable scale. Furthermore, the company competes with open-source frameworks, internal data science teams, and hyperscaler cloud platforms. In Q1, the company debuted its new AI platform that integrates [Microsoft](#) Azure [OpenAI](#) service for data science code generation, yet focuses on deploying its own models.

6: “[How DataRobot Hit \\$210M Revenue with 850 Customers in 2022](#),” GetLatka, 2022, accessed June 15, 2023.

Leadership

[DataRobot](#) has turned over its leadership team as a result of missed revenue results and abuse of stock-based compensation plans. Employees complained of premature stock sales, cronyism, and bloated administrative spending, pushing then-CEO Dan Wright out the door early in Q3 2022. The CFO, Chief Go-To-Market Officer, Chief Product Officer, and Chief People Officer also left. After Wright’s departure, the company promoted President and COO Debanjan Saha to CEO. Saha formerly managed sales and operations for [AWS](#) cloud services and [Google](#)’s data analytics, preparing him to address customer needs. We believe the company will shift its product strategy from providing a universal platform to providing a suite of cloud-based services under Saha’s leadership. The company may also achieve better relationships with cloud providers and superior go-to-market motions. A search continues for a permanent successor to Wright.

Competitors

Due to its breadth of features, as well as its best-in-class augmented analytics for nonspecialist users, [DataRobot](#) has emerged as a leader in the burgeoning autoML market. Smaller competitors



SELECT COMPANY HIGHLIGHTS: DATAROBOT

such as H2O.ai lack the MLOps functionality to manage model deployment. Larger competitors, including [Alteryx](#), [Microsoft](#), and [Google](#), lack the flexibility to customize models across open-source frameworks and production environments. Relative to hyperscalers, [DataRobot](#) offers a superior user experience and support for diverse organizational use cases.


Outlook

[DataRobot](#) must make strategic choices after missing revenue targets and losing leadership due to stock-based compensation scandals. The company’s platform faces competitive pressure from

multiple directions and enterprise customers that are wary of shifting all their valuable data into a separate cloud platform. The valuable components of the platform, including data annotation, MLOps, model search, compliance automation, and model explainability, can be unbundled and integrated with existing systems of record. An open ecosystem approach can build stronger partnerships within the ecosystem and lead to faster sales cycles with opportunities for upselling. We believe that hyperscaler and legacy AI platform partners can benefit from the user experience innovations that [DataRobot](#) has brought to market and the leading MLOps assets the company has acquired in [Algorithmia](#) and Zeff. As it’s currently positioned, the company may be a more likely takeover candidate for PE than an independently large company.

Financing history

Series A	Series B	Series C	Series D	Series E	Series F
August 1, 2014	June 1, 2016	November 3, 2017	October 25, 2018	September 17, 2019	December 9, 2020
Total raised: \$21.0M	Total raised: \$33.0M	Total raised: \$68.0M	Total raised: \$100.0M	Total raised: \$206.0M	Total raised: \$371.0M
Pre-money valuation: \$50.0M	Pre-money valuation: \$89.0M	Pre-money valuation: \$215.0M	Pre-money valuation: \$625.0M	Pre-money valuation: \$1.1B	Pre-money valuation: \$2.5B
Investor(s): NEA	Investor(s): NEA	Investor(s): NEA	Investor(s): Meritech Capital Partners, Sapphire Ventures	Investor(s): Sapphire Ventures	Investor(s): Altimeter Capital Management



About PitchBook Emerging Tech Research

Independent, objective, and timely market intel

As the private markets continue to grow in complexity and competition, it's essential for investors to understand the industries, sectors and companies driving the asset class.

Our Emerging Tech Research provides detailed analysis of nascent tech sectors so you can better navigate the changing markets you operate in—and pursue new opportunities with confidence.

©2023 by PitchBook Data, Inc. All rights reserved. No part of this publication may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, and information storage and retrieval systems—without the express written permission of PitchBook Data, Inc. Contents are based on information from sources believed to be reliable, but accuracy and completeness cannot be guaranteed. Nothing herein should be construed as any past, current or future recommendation to buy or sell any security or an offer to sell, or a solicitation of an offer to buy any security. This material does not purport to contain all of the information that a prospective investor may wish to consider and is not to be relied upon as such or used in substitution for the exercise of independent judgment.

PitchBook Data, Inc.

John Gabbert Founder, CEO

Nizar Tarhuni Vice President, Institutional Research and Editorial

Paul Condra Head of Emerging Technology Research

Additional research

Eric Bellomo
eric.bellomo@pitchbook.com
Gaming
E-Commerce

Brendan Burke
brendan.burke@pitchbook.com
Internet of Things
Information Security
Artificial Intelligence & Machine Learning

Aaron DeGagne
aaron.degagne@pitchbook.com
Medtech
Digital Health

Alex Frederick
alex.frederick@pitchbook.com
Agtech
Foodtech

Jonathan Geurkink
jonathan.geurkink@pitchbook.com
Supply Chain Tech
Mobility Tech

Derek Hernandez
derek.hernandez@pitchbook.com
Enterprise SaaS

Ali Javaheri
ali.javaheri@pitchbook.com
Emerging Spaces

Robert Le
robert.le@pitchbook.com
Insurtech
Crypto

John MacDonagh
john.macdonagh@pitchbook.com
Carbon & Emissions Tech
Clean Energy Tech

Rebecca Springer
rebecca.springer@pitchbook.com
Healthcare Services
Healthcare IT

Rudy Yang
rudy.yang@pitchbook.com
Enterprise Fintech
Retail Fintech

Coming soon
Biotech
Pharmatech