


EMERGING TECH RESEARCH

Artificial Intelligence & Machine Learning Overview

Industry and taxonomy update with latest VC activity

2023

Published April 10, 2023





Contents

Vertical overview	3
VC activity	5
AI & ML landscape	6
AI & ML VC ecosystem market map	7
Segment overview	11
Horizontal platforms	12
Vertical applications	17
Autonomous machines	23
AI & ML semiconductors	26
Appendix	31

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 **Julia Midkiff** and **Joey Schaffer**

Published on April 10, 2023



This report is an annual overview of the AI & ML industry. For a more granular perspective on the AI & ML industry, please see our latest quarterly edition.



Vertical overview


Foundation models present the future of the AI VC ecosystem, as evidenced by recent breakthroughs. [OpenAI](#), [Anthropic](#), [AI21 Labs](#), [Stability AI](#), [Midjourney](#), and [Cohere](#) have achieved groundbreaking AI training results using VC funding to drive the AI field further than tech giants have managed to. Startups have benefited from willingness to take risks on large training runs, utilization of unproven architectures, and agility in product design in commercializing cutting-edge models. Breakthroughs in large language models can produce new business models around mass market interaction with AI systems, including premium services that developers are willing to pay for on a volumetric basis. The availability of intelligent foundation models alleviates the fixed costs of building new AI systems that data scientists have faced in the past and allows scalable consumption of advanced analytics.

Public markets are rewarding enterprises for taking the lead in AI, which improves our outlook for AI adoption. Along with consumers, enterprises also raced to incorporate ChatGPT into their products, leading to significant share price gains for downtrodden names, including [Buzzfeed](#), [C3 AI](#), [SoundHound](#), [Veritone](#), and [BigBear.ai](#), upon announcements of new AI capabilities. This momentum has resulted in large businesses in AI cloud services and model development platforms for hyperscalers. Even so, adoption has remained too uneven across categories to support high-growth startups that require high net retention within customers. Deployments can remain limited after long sales cycles, thereby decreasing the long-term value of a new customer.

The outlook for adoption can drive larger startup businesses over the next three years. Recent data suggests that widespread deployments within organizations are piecemeal yet are reaching critical mass: A survey of chief information officers (CIOs) conducted in June 2022 suggests that most organizations have widespread AI deployments throughout their IT and finance departments.¹

1: "CIO Vision 2025: Bridging the Gap Between BI and AI," MIT Technology Review Insights, Denis McCauley, September 2022.

CIOs' current and expected rate of widespread AI usage by core business function



Source: Data reproduced with permission from [MIT Technology Review Insights](#) | Geography: Global
*As of September 20, 2022




VERTICAL OVERVIEW

This indicates that individual functions can support contract values, if not an entire organization. Looking forward, CIOs plan to infuse AI across departments by 2025, enabling sales cycles to decrease for application-focused startups. 60% of CIOs plan for widespread AI deployment across enterprise functions by then. The ability for large language models to drive enterprise value across functions can finally yield a definitive AI-native enterprise software company that can compete with incumbents such as [Microsoft](#), [IBM](#), and [SAP](#).

We estimate the AI & machine learning (ML) market reached \$197.5 billion in end user spending in 2022, led by the vertical applications segment, along with significant contributions from semiconductors and autonomous machines. AI market research data is improving, given increased research into the attach rates of genuine AI technology to software products. Few AI applications are achieving hypergrowth near 100%, yet many are growing 20% to 30% with the potential to continue growing at a high rate over the next three years. AI-centric software has accelerated past hardware as demand for computing power slows and more applications are constructed atop already-built data centers. AI core software remains a relatively smaller market at \$18.9 billion, excluding natural language processing (NLP) and computer vision applications that overlap with vertical applications. This estimate demonstrates that significant spending will occur in both the platform and application layers of AI, with hardware increasingly becoming commoditized as computing costs decline.

AI & ML market size estimate (\$B) by segment*



Source: PitchBook Emerging Tech Research | Geography: Global | *As of December 31, 2022



AI & ML landscape

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

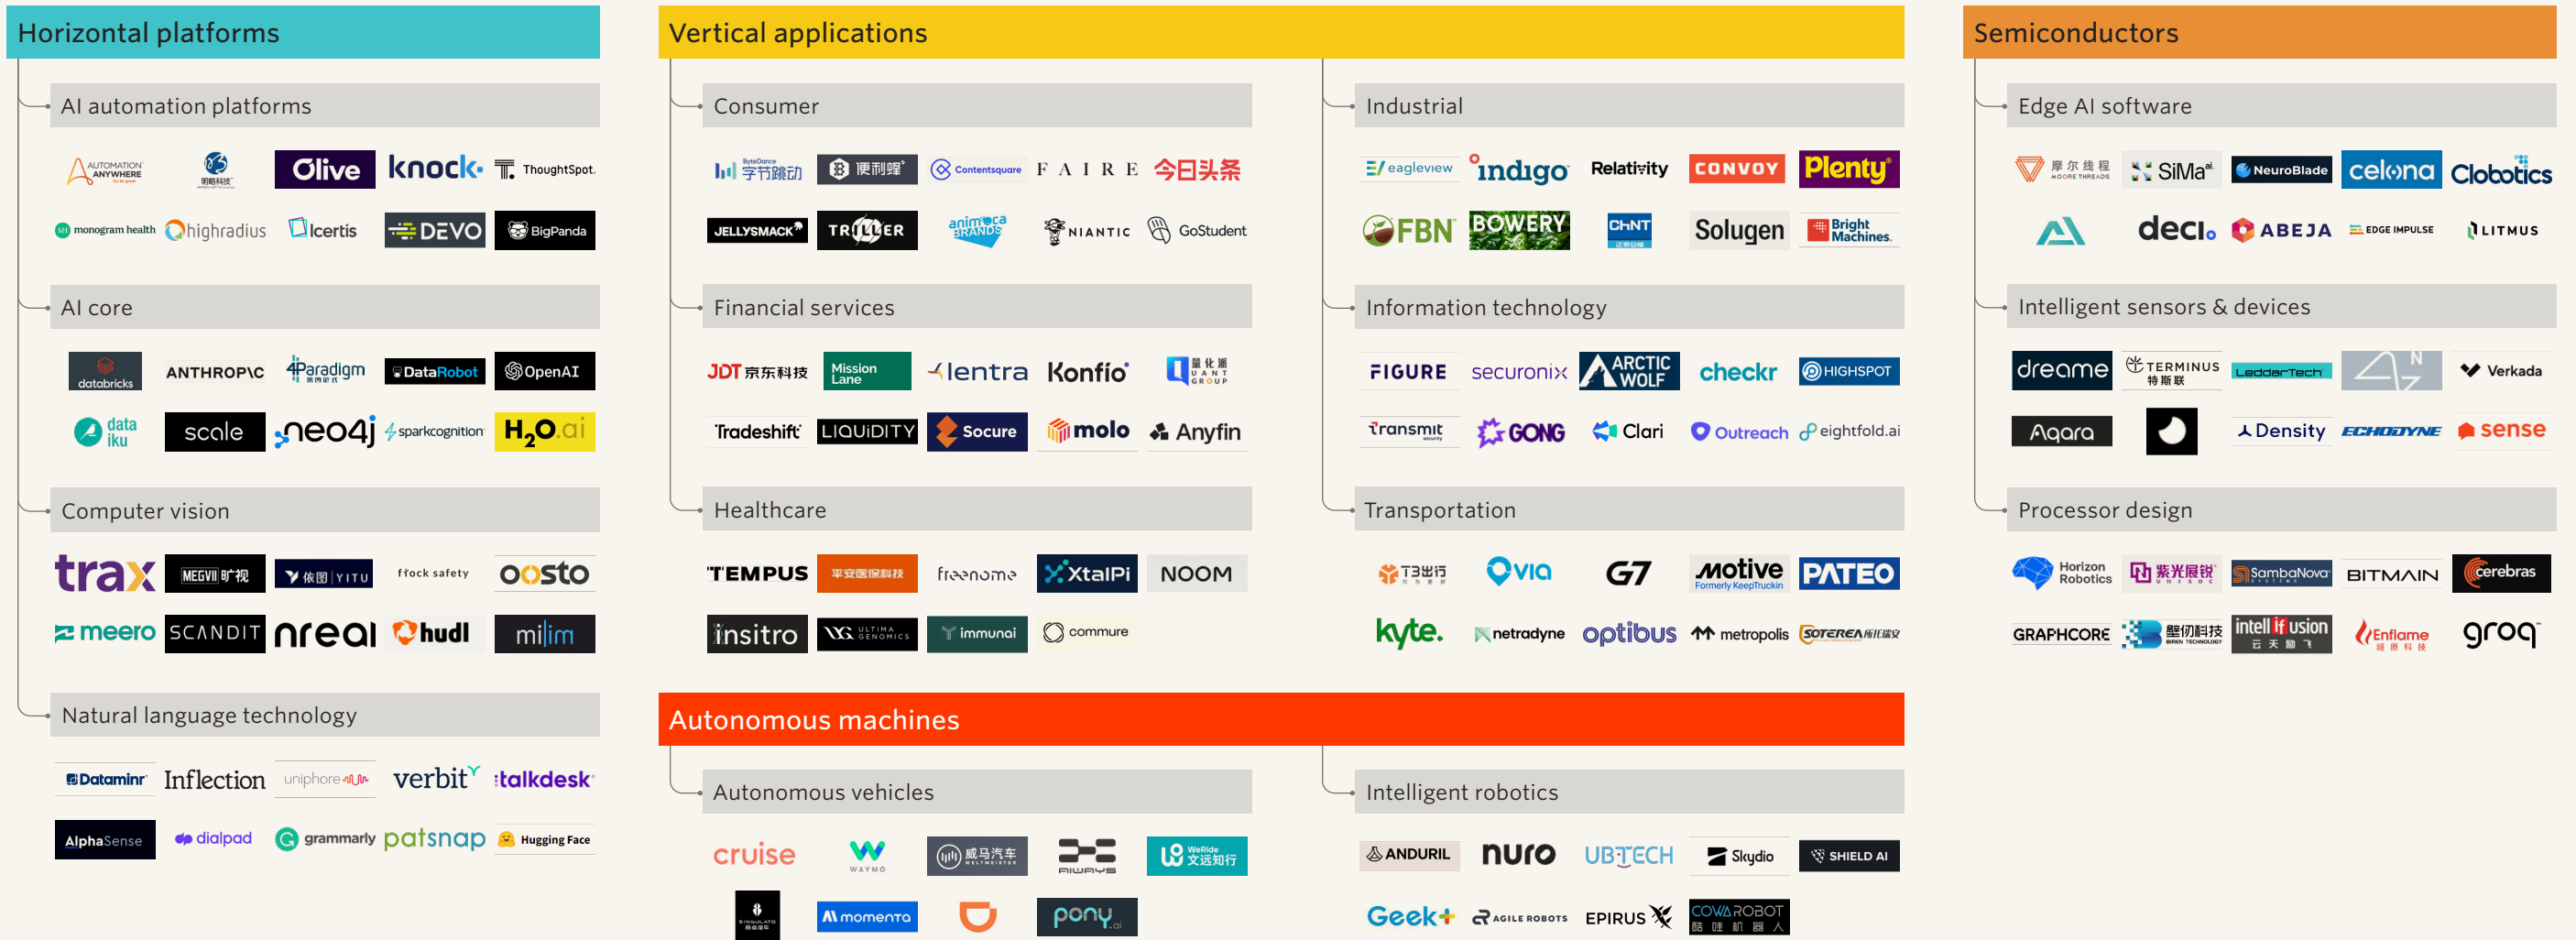




AI & ML VC ecosystem market map

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.






VC activity

2022 challenged the highly funded AI vendors that are unlikely to drive the field's future. VC funding for the vertical declined more than for IT overall, falling 34.9% in VC deal value to \$78.0 billion. The median venture growth valuation step-up declined below the overall VC median, to 1.3x, demonstrating the prevalence of flat rounds and outlying down rounds faced by leading vendors such as [Dataiku](#) and [Noodle.ai](#). As with global IT, 2022 deal value exceeded 2020's total. Deal count also declined more than for IT, at 13.2%, as early-stage deal count regressed even as late-stage deal count remained high. The sector experienced deal count declines of over 50% in autonomous vehicles, processor design, automation platforms, and intelligent sensors as investors avoided hardware- and labor-intensive businesses. Consumer AI proved to be the most resilient vertical application, declining only 20.5% in VC funding year-over-year, as consumer interest in AI applications emerged.

VC exit count continued to decline in Q4, though deal value rebounded due to significant listings in China. We tracked \$7.0 billion in VC exit value in Q4—up from Q3, but the second-lowest value we have tracked since Q2 2020. Exit count has remained above pre-pandemic levels. Among tech giants, active acquirers included [Meta](#), [Oracle](#), and [Spotify](#). PE found opportunities in generative AI, with growth equity PSG using generative copywriting startup Copywriter AI as a platform to add on two writing startups in [Frase](#) and [Rytr](#), forming a new platform called Copyrytr. This platform play gives an early look at the synergies of generative AI startups as they proliferate in different use cases. AI copywriting leader [Jasper](#) has also become an acquirer as part of the company's strategy to deliver a unified content suite across formats.


Median AI & ML VC pre-money valuation step-up by stage




Source: PitchBook | Geography: Global | *As of December 31, 2022



VC ACTIVITY



Source: PitchBook | Geography: Global | *As of December 31, 2022




Source: PitchBook | Geography: Global | *As of December 31, 2022




VC ACTIVITY

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




Source: PitchBook | Geography: Global | *As of December 31, 2022

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




Source: PitchBook | Geography: Global | *As of December 31, 2022

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



Source: PitchBook | Geography: Global | *As of December 31, 2022

AI & ML VC exit count by type



Source: PitchBook | Geography: Global | *As of December 31, 2022



VC ACTIVITY

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

Company	Close date (2022)	Category	Exit value (\$M)*	Exit type	Acquirer(s)/index	MOIC
Almotive	December 22	Autonomous vehicle software	N/A	Acquisition	Stellantis	N/A
Helpshift	December 7	Conversational AI	\$75.0	Acquisition	Keywords Studios	1.4x
Reposify	November 29	Information security automation	N/A	Acquisition	CrowdStrike Holdings	N/A
Tempo	November 2	Processor design	\$689.0	Public listing	N/A	N/A
ECARX	November 18	Autonomous vehicle design	\$3,400.0	Public listing	COVA Acquisition	N/A
Audio Analytic	October 31	Natural language understanding	N/A	Acquisition	Meta Platforms	N/A
DataJoy	October 19	Sales & marketing	N/A	Acquisition	Databricks	N/A
Frase	October 6	Media & entertainment	N/A	Buyout	Copysmith Artificial Intelligence, Harmony Venture Labs, PSG	N/A
Rytr	October 6	Media & entertainment	N/A	Buyout	Copysmith Artificial Intelligence, Harmony Venture Labs, PSG	N/A
Blue Hexagon	October 4	Information security automation	\$10.0	Acquisition	Qualys	0.3x

Source: PitchBook | Geography: Global | *As of December 31, 2022



Segment overview

Horizontal platforms

Open-source adoption accelerates commercialization of generative AI for the mass market.

AI & ML semiconductors

Struggling data center startups encourage innovation in edge and photonics applications.

Vertical applications

Opportunities in revenue operations, defense, and media stand out.

Autonomous machines

Self-driving software startups respond to declining funding with innovation and practical goals.



Horizontal platforms

Overview

Horizontal platforms empower end users to build and deploy AI & ML algorithms across a variety of use cases. These platforms directly apply scientific advances in AI & ML research to commercial applications. Companies in this segment have differentiated AI & ML approaches and are built with AI & ML from the ground up—this is also referred to as AI-first. Furthermore, some horizontal platforms are used to improve AI & ML algorithms but do not use AI & ML themselves.

Subsegments include:

- **AI core:** Building blocks of AI & ML deployments, including developer tools needed to build and deploy models to production. Categories within this subsegment include AI as a service (AaaS), AI & ML developer tools, AI platform as a service (PaaS), automated ML (autoML), cognitive computing, data preparation platforms, quantum AI, and TinyML.
- **Computer vision:** The use of AI & ML to analyze visual data and make meaningful predictions about both the physical world and digital images. The technology can be used across use cases to label and make predictions about visual data. Key products utilizing computer vision across a range of verticals include AI-enabled augmented reality, computer vision as a service, facial recognition, geospatial analysis, and visual data labeling software.
- **Foundation models:** Foundation model companies train custom neural networks using generalizable techniques including transformers, diffusion models, and multimodal approaches. Companies in this space incur substantial computing costs to compete on academic benchmarks and commonly have artificial general intelligence (AGI) as a long-term goal.

- **Natural language technology (NLT):** NLT uses computational linguistic techniques to learn from communications data and make predictions about the structure and content of language. Categories within this subsegment include conversational AI, neural machine translation, natural language generation, NLP, and natural language understanding.
- **AI automation platforms:** Software and services that enable enterprises across all verticals to leverage AI to automate critical business processes via predictive analytics. Categories include AI-first applications of the following products: AI for IT operations (AIOps), business intelligence, contract lifecycle management automation, database management, decision intelligence, and intelligent process automation.

Industry drivers

Costs of foundation model training coming down: GPT-3 was highly costly to train as of its launch in 2020, with estimates ranging around \$10 million for a single training run. Stable Diffusion brought the cost for a state-of-the-art generative model down to around \$600,000.² Hardware optimization can produce further cost savings, with competitive foundation models able to be trained at a 5x to 20x lower cost using customized hardware and acceleration software from vendors such as [SambaNova](#) and [MosaicML](#). As an example of this, [AI21 Labs](#) was able to train a competitive large language model after raising only \$35.0 million in VC funding.

Talent shortage: Hiring AI data scientists has become more difficult over the past three years for 47% of AI leaders, according to a global survey. 78% of organizations find hiring AI data scientists difficult. Buying AI software can address this talent shortage.³

²: [“We actually used 256 A100s for this per the model card, 150k hours in total so at market price \\$600k,”](#) Twitter, Emad Mostaque, @Emostaque, August 28, 2022.

³: [“The State of AI in 2022—and a Half Decade in Review,”](#) McKinsey & Company, December 6, 2022.



HORIZONTAL PLATFORMS

Market size

We estimate that spending on this segment’s categories reached \$32.4 billion in 2022 and will grow at a 32.7% CAGR to \$75.7 billion in 2024. This market size includes our estimate of spending on computer vision and NLT applications—even when they overlap with vertical applications spending. We estimate that AI core software, including AI data preparation, model training, and infrastructure platforms, amounted to \$14.0 billion including AI platforms and model development tools. More precise data collection for horizontal computer vision software resulted in an estimate of \$1.8 billion, below our previous estimate that overlapped with vertical applications. AI automation platforms, which include AI-integrated business analytics and robotic process automation (RPA) platforms, reached \$4.9 billion in 2022.


Business model

Horizontal platforms typically contain a variety of computational costs, including the following components for a typical hyperscaler ML deployment:

Infrastructure costs:

- **Storage:** Data is typically stored in a cloud server and carries nominal cost based on the scale of the data.
- **Network:** Data moved into and out of the cloud environment is charged on a volume-based rate.
- **Compute:** Model building includes separate costs for building, training, and deployment.

Horizontal platforms market size estimate (\$B)*



Source: PitchBook Emerging Tech Research | Geography: Global | *As of March 15, 2023



HORIZONTAL PLATFORMS

Operational costs: Cloud servers carry an hourly rate that can vary by the contract length and compute requirements, among other factors.

Security & compliance costs: AI & ML data must have an audit trail for various data privacy compliance frameworks.

Each of these components has separate costs for both training and deployment. As a result, we estimate training a single model of GPT-3's scale can cost around \$2.5 million in compute resources. Optional incremental costs include breach discovery and log management, maintenance and support costs, data encryption, and secure web gateways. Startups can unbundle hyperscaler platforms and provide data preprocessing and monitoring on top of the hyperscalers' platform.

Opportunities

Generative AI: Open-source adoption accelerates commercialization of generative AI for the mass market. In its first three months after launching, [Stability AI's](#) Stable Diffusion Model accrued nearly 34,000 stars on [GitHub](#), a leading measure of popularity among developers. [Hugging Face](#) brought language models to the open-source community and has since reached a \$2.0 billion valuation in its Series C only six years after founding. Open-source models' ability to create new commercial applications can be seen by the derivative products using Stable Diffusion, including [Playground AI](#), [Google Research's](#) DreamBooth, [Replicate](#), and RunwayML. Stable Diffusion has been forked over 5,000 times since its launch on [GitHub](#), showing the number of developers

building new projects and commercial products on top of the software. The growth curve for open-source AI has been redefined and shows that generative AI techniques can spur developer excitement in a range of fields.

At the early stage, startups demonstrate that foundation models can support a range of business models. While [Jasper](#) has shown that [OpenAI's](#) large language model GPT-3 can power a marketing copywriting business, [Regie.ai](#) exhibits similar growth by using the foundation model to standardize sales content across channels. In addition to the contextual understanding, the company has retrained models on 35 million sales emails, giving the company an exclusive advantage in a tangential function. The startup achieved a 4.3x valuation step-up into its Series A, resulting in a \$55.0 million post-money valuation. Emerging opportunities exist in video synthesis and editing, as evidenced by [Peech's](#) \$8.3 million seed round, as well as audio generation. For more on the generative AI opportunity, see our [2023 Vertical Snapshot: Generative AI](#).

Data-centric AI: Even with the success of foundation models, data quality remains paramount to model accuracy. [OpenAI's](#) ChatGPT relied on manual data labeling from startup [Sama](#) to complement the raw outputs of the company's large language model. [Sama's](#) data labelers scored ChatGPT's responses to validate accurate responses, which resulted in a highly reliable chatbot with capabilities across question answering, software development, and text generation. [Scale AI](#) continues to take a leadership role in the AI community, given its competitive advantage in data labeling for computer vision use cases. The company achieved a \$7.3 billion post-money valuation in 2021.



HORIZONTAL PLATFORMS

Given the market downturn, mega-deals did not occur in this space during 2022, yet late-stage vendors demonstrated continued growth. After launching in 2021, data engineering startup [Galileo](#) raised an \$18.0 million Series A to automate correction of data errors. The company was founded by former ML engineers from [Google AI](#), [Uber](#), Stanford's AI Lab, and Carnegie Mellon. Other significant late-stage rounds included Series B rounds from [DataLoop](#) and [Superb AI](#) that valued the companies over \$100 million. We believe large companies can be created in this niche.

Streaming database management: Database management innovators are focusing on stream processing as part of their AI strategies. Both [Snowflake](#) and [Databricks](#) announced new streaming products at their conferences in June 2022. [Databricks](#) notes that streaming jobs with its Structured Streaming platform doubled from around 2 million to over 4 million YoY as of July 2022 after launching in January 2019.⁴ At its Data + AI Summit, the company launched Project Lightspeed to lower the latency of its Structured Streaming product, particularly adding support for the Python application programming interface (API), which we believe will better facilitate machine learning inference on streaming data. At [Snowflake's](#) Summit, the company announced Snowpipe, a streaming data ingestion service. These vendors are responding to customer demand and the long-term need for continuous AI inference.

Data integration startups with support for stream data pipelines are achieving high acquisition interest. In Q2, [Software AG](#) closed its acquisition of StreamSets for \$584.0 million, granting the company a multiple on invested capital (MOIC) of 7.6x. The acquisition was justified by StreamSets opening a total addressable market of \$3.5 billion in data integration by 2025. Previously, in Q3 2021, [Fivetran](#) acquired data integration startup [HVR Software](#) for \$700.0 million, demonstrating the high exit values that can be achieved in this space. In VC, Astronomer achieved a 5.9x

4: ["Project Lightspeed: Faster and Simpler Stream Processing With Apache Spark," Databricks, June 2022.](#)

valuation step-up and unicorn status in its Series C to orchestrate data pipelines. Legacy database companies can benefit from partnering with startups to integrate real-time data from new sources including cloud, Internet of Things (IoT), and web applications.

Risks & considerations

Public companies struggling to achieve widespread AI adoption: Public AI companies bringing AI to a wide audience are struggling to gain high revenue growth. Pure-play horizontal AI companies that enjoyed high valuation premiums earlier in the COVID-19 pandemic, including [Palantir](#), [C3 AI](#), and [Veritone](#), face challenging market conditions as technology valuations have regressed below pre-pandemic multiples and enterprises shift spending. During the market downturn, public AI companies have become the target of short sellers aiming to expose inflated claims around future contract value. These companies have also faced declining revenue growth and limited expansion beyond their core customer segments.

Public cloud hosts offer comprehensive capabilities for AI & ML model building and deployment: Public cloud hosts can offer AutoML, language, and vision services. They also feature AI marketplaces, including products such as APIs, microservices, datasets, and prebuilt algorithms, that compete with horizontal platforms. Incumbents with offerings in this space include [Alibaba](#), [Amazon](#), [Apple](#), [Baidu](#), [Microsoft](#), [Google](#), [Tencent](#), and [IBM](#). We believe enterprises with limited budgets can spin up limited AI & ML projects with a combination of open-source frameworks and cloud management tools. For this reason, scaling of multicloud and API-based ML operations (MLOps) tools may be delayed until the return on investment (ROI) for existing AI & ML projects is strong enough to encourage heavier investment in AI & ML capabilities across the enterprise.



HORIZONTAL PLATFORMS

AI ethics and regulation: Deep learning is compounding the black-box nature of AI & ML as it becomes more effective, because its features are inherently fluctuating and obscure to even the data scientists training them. Models can lack clarity on critical metrics including privacy, security, ethics, and transparency, thus limiting their utility in sensitive use cases. Active regulatory efforts are ongoing from the US federal government, the EU, China's internet regulator, and the UN High Commissioner for Civil Rights. Explainability is critical to understanding ethical issues in training data bias and discriminatory model parameters but is unavailable without a ground-up focus during the requirements stage. Vendor differentiation will likely emerge based on the level of model transparency offered along with ethical data collection.



Vertical applications

Overview

Vertical applications in AI & ML address specific problems within industries and are not always AI first. Many startups in this category design a solution to an industry problem using software and integrate AI & ML to optimize some part of their product. These solutions typically differentiate based on the quality of the dataset used to train the industry-specific model and the industry expertise of the data scientists identifying decision-making areas that can be enhanced by AI & ML models. As a result, many of these startups help automate specific functions within their industry but have limited ability to cross apply their AI & ML to other industries.

Current subsegments include:

- **AI & ML in financial services:** Includes technologies that embed AI & ML into existing financial services via advanced analytics, process automation, robo-advisors, and self-learning programs. Product categories include financial chatbots, intelligent banking, lending analytics, payment optimization, predictive underwriting, and robo-advisors.
- **AI in healthcare:** Includes technologies that leverage AI & ML to improve medicine and the provision of care. Product categories include AI-based drug discovery, clinical decision support, genetic analytics, healthcare administration, and personal health.
- **Consumer AI:** Includes technologies that use AI & ML to enhance B2C business models. Product categories include AI in media & entertainment, AI & ML advertising technology, digital avatars and gaming, e-commerce recommendation engines, education technology, intelligent price optimization, and smart retail.
- **Industrial AI & ML:** Includes technologies that automate industrial processes and unlock industrial data to find new efficiencies. Product categories include crop maximization, energy

grid automation, geospatial analysis, heavy industry automation, IoT predictive analytics, supply chain optimization, and telecommunications optimization.

- **AI in IT:** Includes enterprise software tools that optimize specific functions typically administered by IT departments, including both back-end and front-end use cases. Product categories include human resource automation, information security (infosec) automation, IT infrastructure management, legal automation, property technology, sales & marketing automation, and software development tools.
- **AI in transportation:** This category includes intelligent fleet management and driver assistance. It excludes autonomous vehicles.

Industry drivers

Growing availability of industrial datasets: Startups can build preliminary models based on data and frameworks from sources including [Google Dataset Search](#), [Kaggle](#), UCI Machine Learning Repository, [Microsoft Coco](#), and [GitHub](#). Governments support data sharing, as evidenced by President Biden's Executive Order on Enhancing Safeguards for United States Signals Intelligence Activities.⁵ This order codifies the EU-US Data Privacy Framework for secure data sharing for AI training.

Enterprise customers are achieving ROI with AI solutions: McKinsey's State of AI survey found that over 40% of AI adopters in service operations, manufacturing, human resources, risk management, and sales & marketing achieved cost savings as a result of their AI projects in 2021.⁶ Only around 50% of these adopters reported cost savings the prior year.

5: "[Executive Order on Enhancing Safeguards for United States Signals Intelligence Activities](#)," The White House, October 7, 2022.

6: "[The State of AI in 2022—and a Half Decade in Review](#)," McKinsey Institute, December 6, 2022.



VERTICAL APPLICATIONS

Market size


We forecast that the vertical applications market reached \$94.0 billion in 2022, with a 31.2% CAGR out to 2025, resulting in a \$212.2 billion market. The segment is weighted heavily toward enterprise IT applications, which we view as a \$28.6 billion market, led by sales, infosec, and human resources. We forecast IT to grow to a \$65.4 billion market by 2025 at a 31.7% CAGR. We estimate industrial AI to be the second-biggest category, at \$16.3 billion, focused primarily on manufacturing automation, supply chain optimization, and predictive maintenance. Both the consumer AI and AI in financial services markets exceed \$10 billion. Healthcare is a smaller market, at \$4.1 billion, given barriers to adoption for third-party AI solutions. This estimate has been revised upward due to new data on spending for digital assistants and media augmentation.

Business model

The business models of enterprise AI vendors can vary widely and have considerable influence on their cash flow profiles. Software AI startups employ several key business models, including:

- **SaaS:** A monthly subscription that bundles data engineering, model building, and deployment. Additional setup and maintenance fees may be layered on top of a base package. Applications can be delivered as a cloud service through channel partners. In most cases, managed services are required via the vendor or a third-party consulting firm.
- **Codevelopment agreements:** An AI-first vendor partners with an industry specialist to codevelop an AI solution to be used by the industry specialist as a customer and licensed to the industry specialist for distribution. Often, this includes an upfront payment to the AI-first vendor

Vertical applications market size estimate (\$B)*



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



VERTICAL APPLICATIONS

to develop the initial models. This arrangement can be necessary to supply AI-first startups with the requisite data to enter a new industry. It is a low-margin business model compared with intellectual property (IP) ownership.

- **Joint venture:** An AI-first vendor and enterprise create a joint venture to develop vertical applications and market them. Typically, the customer is given the choice to buy out the vendor at a future date.

In each case, AI-first vendors bear costs related to data engineering, model building, and maintenance, thereby creating a cost structure that exceeds those of SaaS business models. Consulting firms also arrange outcome-based or upfront capital expenses, though these are not commonly employed by startups. In each case, startups must require customers to consent to data contribution to improve their models without a royalty fee for the customer.

Opportunities

Code completion: AI code completion has begun to take off after years of false starts. Recent innovations in transformer models make language models into effective coders, having learned from open-source code samples. We estimate the intelligent software assistant market will reach \$2.7 billion in spending in 2023, offering a commercial opportunity for startups. Startups in the software development tools category have not produced a unicorn valuation yet, although [Builder.ai](#) reached an \$815.0 million post-money valuation in its February Series C. After commercial successes applying GPT-3's natural language understanding, innovative startups have already integrated [OpenAI's](#) Codex in commercial products including [Replit](#), [Pygma](#), and [Machinet](#). AI will be used for limited code completion over the next two years, contributing a minority of production codebases in this first phase of commercialization.

At the earlier stage, [TabNine](#) raised a \$15.5 million Series B led by strategic investors to pursue the code completion vision and compete with [GitHub's](#) Copilot product. The company has been developing predictive models on a range of programming languages since 2012. It develops customized models for each programming language, compared with the monolithic Codex model that forms the basis of [GitHub](#) Copilot. The company has not achieved significant commercial scale, as its product may have arrived in the market too early. The recently improved performance of language models presents a catalyst for this company, given improved accuracy of code prediction.

Revenue operations intelligence: Increasingly, enterprises are creating a single department called revenue operations (RevOps), which integrates sales, marketing, and customer success with a single IT back-end and set of common industry key performance indicators (KPIs). Customers can connect various data sources including customer relationship management (CRM), social media, enterprise resource planning (ERP), and third-party analytics platforms into a central platform that uses ML to normalize disparate data streams and uncover correlations among them. Market leaders in RevOps include AI-native startups such as [Gong](#), [Outreach](#), and [Clari](#). [Gong](#) gained a foothold with conversational AI for sales calls and is expanding into forecasting and recommendations based on its unique insights. In Q1 2022, the company released a direct integration with [Salesforce](#) to enrich data in Sales Cloud. [Clari](#) and [Outreach](#) offer more direct competition with [Salesforce](#) by offering comprehensive revenue management solutions that natively integrate capabilities across forecasting, conversation intelligence, and pipeline outreach automation.

Sales automation stands out as a leading use case of generative AI. While [Jasper](#) has shown that [OpenAI's](#) large language model GPT-3 can power a marketing copywriting business,



VERTICAL APPLICATIONS

[Regie.ai](#) exhibits similar growth by using the foundation model to standardize sales content across channels. In addition to the contextual understanding, the company has retrained models on 35 million sales emails, giving the company an exclusive advantage in a tangential function. The startup achieved a 4.3x valuation step-up into its Series A, resulting in a \$55.0 million post-money valuation.

Sales & marketing automation has separated itself as the most active category of VC exits in AI. The category produced a segment-leading 20 VC exits in 2022 after also leading in 2021 with 34. While we did not track a high volume of large exits in the down market of 2022, significant acquirers remained active, including [LinkedIn \(Oribi\)](#), [Salesforce \(Troops\)](#), and [Databricks \(Datajoy\)](#). [Troops](#) converts AI analytics into timely messages for Slack, thus demonstrating the value of generative content for sales automation. We expect momentum around generative applications to encourage further M&A and create some independently large companies.

National defense: The US Department of Defense's (DoD) Chief Digital and Artificial Intelligence Officer, Dr. Craig Martell, is requesting \$615.2 million for demonstration and validation of new technologies in FY 2024—more than double FY 2023's budget.⁷ Given the strategic priority of competition with adversaries including China, we expect total DoD spending on AI to reach \$1.8 billion by 2024. In 2022, the Department of Defense awarded 79 contracts to artificial intelligence vendors, including 44 small businesses. The success of small businesses demonstrates the Pentagon's streamlined procurement process.

Defense startups are achieving some of the greatest scale of any AI startup through the market downturn, given global geopolitical challenges. Autonomous warfighting startup [Anduril](#) achieved

an \$8.5 billion post-money valuation in a Series E. The company won a 10-year contract for nearly \$1 billion to assist the Department of Defense with drone defense using a customized autonomized drone to ram into enemy drones. Pure-play AI startup [Scale AI](#) won a \$250 million data labeling contract in 2022, demonstrating the scale that can be achieved by startups working in defense.

Metaverse: We believe that AI & ML has an important role to play in the creation of a Metaverse, which we define as a virtual world in which gaming and social media can come together to create novel interpersonal and entertainment experiences. [Meta's](#) AI Research Supercluster—now one of the fastest supercomputers in the world—is designed to train large language models, which the company views as core to its Metaverse ambitions. The compute cluster promises to enable real-time machine translation for multilingual conversations in the Metaverse while also detecting abusive content. AI & ML is also necessary for the computer vision behind realistic avatars. Gaming giant [Roblox](#) is actively pursuing generative AI models including Stable Diffusion to add content more rapidly than human teams can on their own.

Startups are seizing the Metaverse opportunity with both content generation and conversational moderation. On the generation side, [Klang Games](#) reached a \$403.2 million post-money valuation in its Series C led by AI gaming innovator [Animoca Brands](#). The company uses AI to generate characters for a large-scale multiplayer game that persists even once players log off. In moderation, [Modulate](#) achieved a \$170.0 million post-money valuation in its Series A to use AI to monitor abusive speech in virtual environments. We believe large businesses can be created in both content creation and moderation and expect to see M&A as generative AI creates new experiences.

⁷: "[DoD Budget Request](#)," Under Secretary of Defense (Comptroller), March 15, 2023.



VERTICAL APPLICATIONS

Risks & considerations

Excessive fixed costs: We believe AI companies commonly spend 25% or more of their revenue on cloud resources. State-of-the-art generative models can cost at least \$500,000 to train, based on current graphics processing unit (GPU) pricing, and are often retrained on a regular basis. Additionally, images, audio, or video data incur storage costs. Lastly, transferring trained models across cloud environments incurs transfer costs, which can flow to costs of goods sold, if model training is required for each customer engagement. Investors should question how startups utilize cloud resources in their model training and inference.

Overfit models: With large datasets, a model can be trained extensively to the point where it has too many features to make predictions on novel datasets. As a result, some models trained on quality datasets may not have “model-market fit,” even if the software solution itself appeals to customers. Startups should engage in early stopping for neural network training and utilize small datasets to avoid this problem.

Biased data: Models trained on racially biased data can contribute to racially biased decisions in financial services, HR automation, healthcare, transportation optimization, and consumer AI. Racial bias in models can emerge from unexpected sources, such as healthcare spending data,

due to sampling bias and unforeseen features including correlating health spending to health risk. For startups, the public relations consequences of a biased model are material, and racially biased models must be retrained on fresh data with additional feature engineering and testing. The effect of such a restart can amount to millions of dollars. Thus, it can be cost effective to integrate explainability and ethical AI governance frameworks from an early stage to ensure that AI algorithms do not incorporate racial bias in their features or predictions. The negative consequences of bias will gain more public attention as AI & ML is increasingly used in lending and hiring decisions.

Privacy concerns: The retail and advertising industries are another AI & ML end market that is rooted heavily in the use of personal data. With recent lapses in protection and governance of personal data, more concerns will likely emerge surrounding data science products in consumer-centric use cases. Regulatory changes, such as the EU’s implementation of the General Data Protection Regulation (GDPR), have far-reaching effects on how AI & ML-focused companies need to operate, especially if other jurisdictions follow suit. Adherence to these higher standards will likely add operational costs and could slow implementation in some geographies. Some AI developers are facing lawsuits for their infringement on data sovereignty. This can also affect individual behavior without the intervention of the government, as potential customers may choose an alternative because of privacy concerns.



Autonomous machines

Overview

Autonomous machines can perform tasks in human-present environments without explicit human control. These machines synthesize ML, computer vision, and datasets of the physical world such as navigation. The segment requires the design of complex hardware with software “brains” and is therefore valued differently than AI & ML software as well as semiconductors. Current applications of AI in robotics lack “common sense” to learn from their environments. This segment does not include optimization of existing robots, including those used in manufacturing and supply chains.

Subsegments include:

- **Autonomous vehicle software and design:** Software and hardware solutions that enable self-driving or driver assistance capabilities for cars, trucks, and other on-road vehicles.
- **Intelligent robotics & drone design:** Robotic systems and uncrewed aerial vehicles that can operate without human input. AI & ML can be used for the learning, control, and adaptation of robots. Common applications of autonomous robotics include swimming, as well as picking up, carrying, and putting down objects.
- **Intelligent robotics & drone software platforms:** Operating systems for autonomous robots. These platforms can include fleet management and predictive maintenance.

Industry drivers

Autonomous delivery demand: Because of the closure of retail outlets, last-mile delivery requires additional capacity with less cost sensitivity from merchants. Companies offering autonomous delivery are stepping in to expand the capacity of overstretched delivery networks.

E-commerce growth driving autonomous robotics implementation: E-commerce has maintained its percentage of retail sales since the pandemic’s lockdowns at around 15%, encouraging automation of warehouses to serve elevated demand.



AUTONOMOUS MACHINES


Market size

We estimate the autonomous machines market to be \$41.1 billion as of 2022 and forecast it to grow at an 18.9% CAGR to \$58.1 billion in 2025. This estimate includes commercial and delivery drones, autonomous vehicles, and autonomous robots. Most of this estimate stems from semi-autonomous passenger car sales. While the industrial robotics market is larger on its own, at over \$100 billion, we focus on the smaller opportunity for autonomous robots within the manufacturing, healthcare, and supply chain industries. Autonomous driving is expected to grow at a 23.3% CAGR through 2024—higher than for robotics, given the challenge of displacing existing industrial robotics in manufacturing and warehouse settings.

Business model

Beyond traditional hardware sales & service business models, autonomous machine suppliers are increasingly delivering robotics as a service to operate with multiple types of hardware. AI & ML algorithms trained for machine intelligence can be deployed on an edge device via software development kits and paid for via recurring revenue. We believe this model is becoming more prevalent among incumbents and is table stakes for startups.

Autonomous machines market size estimate (\$B)*



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



AUTONOMOUS MACHINES

Opportunities

AI for autonomous driving: Vehicle ML algorithms struggle to deal with edge cases including unfamiliar objects and scenarios. The first wave of startups collected much of their data from the perspective of the car's sensors, adding real-world data over time through simulation and testing. Some edge cases, including climate, human behavior, stray objects, and the built environment, have separate datasets that can be used to generate their own models.

Self-driving software startups are continuing to scale through a downturn for self-driving funding. Since we covered the company in our [2020 Annual AI & ML Report](#), [ISEE](#) has evolved to develop an industry-leading solution for logistics yard truck automation. The company has narrowed its vision from a general autonomous trucking solution yet has achieved commercialization, resulting in a \$40.0 million Series B led by Founders Fund with participation from shipping giant [Maersk's](#) corporate VC (CVC) fund. AI-focused startup [Autobrain](#) achieved a 2.7x valuation step-up, which resulted in a \$387.5 million post-money valuation in Q2 2022, showing that self-driving software can benefit from both fully autonomous or semi-autonomous driving. The company is hardware-agnostic.

Warehouse robotics: Warehouse operators are looking for ways to bridge the manual-labor-intensive gap between brick-and-mortar stores and last-mile delivery. [Amazon](#) Mobile Robotics, formerly [Canvas Technology](#), has touted this as a competitive advantage based on its vision-based navigation. Logistics customers are prioritizing robotics investments to respond to the throughput increases and workforce safety needs during the pandemic. Precedent transactions include Ocado Group's acquisition of [Kindred AI](#), [Hyundai's](#) acquisition of [Boston Dynamics](#), and, in 2019,

[Amazon's](#) acquisition of [CANVAS Technology](#). Global warehouse robotics leader [Geek+](#) achieved a \$2.0 billion post-money valuation in Q3 2022, demonstrating the scale of the technology. Recently listed warehouse robotics company [Berkshire Grey](#) is also winning significant contracts, leading its share price to perform well recently.

Autonomous drones: Autonomous drones can unlock new use cases in field monitoring, building monitoring, and extreme environments. We believe construction monitoring is an area of high spending growth, with additional connectivity becoming available in rural areas. Autonomous drones can form a central part of its growth story. We estimate the commercial drone market reached \$3.4 billion in 2022 and forecast it to grow at a 26.9% CAGR to \$6.6 billion in 2025.

Startups [Exyn](#), [Embention](#), [vHive](#), [PINC](#), and [PreNav](#) have sensor-enabled drones that address a range of use cases, from supply chain to subterranean mine inspection. [Exyn](#) has developed autonomous drones with multiple embedded sensors that can enter GPS-denied environments such as subterranean mines. The company raised a \$35.0 million Series B extension at a 1.6x valuation step-up in Q4 2022 from strategic investor Reliance Industries, indicating the commercial potential of the technology.

Risks & considerations

Products killed before commercialization: The lag between product announcements and launch dates can last over a year. During this time, technical roadblocks can emerge, and cost overruns can reduce the economics of already low-margin hardware products. Robotics teams should have extensive experience with go-to-market strategies to overcome these commercialization hurdles.



AUTONOMOUS MACHINES

Robots in human-present environments failing in the field: As new use cases emerge, robotics companies are struggling to turn proofs of concept into sticky revenue. [SoftBank Robotics](#) has faced high churn on its robot-as-a-service contracts, thereby pressuring it to shorten the three-year length of its contracts.

Regulatory risk: In autonomous vehicles and mobile robotics, uncertainty around regulatory frameworks may limit the growth of emerging technologies. The National Highway Traffic Safety Administration currently does not exempt self-driving cars from occupant safety standards, thus limiting the benefits of driverless technology. In robotics, regulation might be reactive to the introduction of robots in human-present environments, which creates risk along with innovation in domain randomization.



AI & ML semiconductors

Overview

AI & ML model training and inference require hardware with maximum computational efficiency and customized processing for AI calculations. AI requires a high volume of contemporaneous calculations that run in parallel and benefit from specialized chips. Using general purpose chips can cost thousands of times more than using an AI chip for these purposes. As a result, a segment has developed for both the design and software-based optimization of computing hardware, including both semiconductors and sensors.

Subsegments include:

- **AI chips:** These types of computer chips attain high efficiency and speed for AI-specific calculations. AI chip optimizations include parallel calculations, low-precision calculations to reduce transistor count, AI & ML algorithm compression, and using novel programming languages built specifically to efficiently translate AI computer code for execution on an AI chip. AI chips are faster and more efficient than central processing units (CPUs) by orders of magnitude for the training and inference of AI algorithms. Technologies within this category include:
 - GPUs
 - Field-programmable gate arrays (FPGAs)
 - AI-specific application-specific integrated circuits (ASICs)
 - Neural network processors

- **Edge AI:** Compression algorithms that optimize AI & ML models for deployment within various semiconductor environments and edge devices. Technologies in this category include:
 - Post-training neural network binarization and quantization algorithms
 - Quantization aware training
 - Ensemble boosting algorithms
 - Edge AI APIs and software development kits (SDKs)
- **Intelligent sensors & devices:** Devices that are optimized to run AI & ML models. This segment contains technologies including:
 - Sensors: Devices that measure specific parameters in real-world conditions such as environmental conditions, motion, images, and chemical levels.
 - Sensor systems: Assemblages of sensors and chipsets that empower specific IoT use cases such as smart buildings, industrial IoT, and connected vehicles.

Industry drivers

Expensive computation: The amount of compute being used in large AI training runs has been increasing exponentially. [NVIDIA](#) estimates that the compute required to train leading transformer models is increasing by 275x every two years.⁸ [Google](#) benefited from its proprietary tensor processing units, but similarly ambitious and practical projects undertaken by other companies may require even higher costs to replicate.

⁸: ["Investor Presentation: Q4 FY23," NVIDIA, February 27, 2023.](#)



AI & ML SEMICONDUCTORS


Research finding that deep learning improves at scale: The more compute that is used, along with optimal model architecture and large data training sets, the more accurate deep learning algorithms become. This result has been empirically demonstrated by ML researchers and means that additional breakthroughs in accuracy will scale in tandem with computing power.

Demand shifting from discrete to integrated processors: AI chips have focused on the data center but are increasingly shifting to the edge to support emerging use cases including smartphones, tablet computers, wearable devices, and connected automobiles. This is leading to integrated processors growing faster than discrete processors.

Market size

We estimate the AI & ML semiconductor market reached \$43.6 billion in 2022 achieving 22.5% growth over 2021. Over half this total can be attributed to mobile phone ASICs, which will decline as a percentage going forward. We expect the market to grow at an 21.8% CAGR from 2022 to 2025, resulting in a \$72.7 billion market. Our estimate includes all AI semiconductor varieties, such as field-programmable gate arrays, GPUs, microcontrollers, and ASICs. To be included in this estimate, these chips' primary function must be enabling AI processing. We forecast the highest growth segments in this market to be field-programmable gate arrays and microcontrollers, although we anticipate all segments will grow at CAGRs over 15.0% from 2022 to 2025, based on emerging edge AI use cases.

AI & ML semiconductor market size estimate (\$B)*



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



AI & ML SEMICONDUCTORS

Business model

AI semiconductors are currently deployed through unit sales and usage-based pricing from cloud hosts. For example, [AWS](#) sells usage-based instances of [NVIDIA](#) GPUs for AI & ML workloads. The usage-based pricing of GPU instances enables around a one-year payback for [AWS](#) for the underlying hardware, assuming constant usage of these GPUs and given current pricing of around \$10,000 for a standard [NVIDIA](#) Tesla A100. AI & ML developers can thus gain access to GPUs without high capex.

Opportunities

Inference semiconductors: AI chips address two critical phases of the model lifecycle: the training of new models based on existing data and the inference of those models on live data in production. Training-focused startups establish benchmarks based on the time required to train and the efficiency of common benchmark models. Inference startups, in comparison, stand out for their speed and efficiency of inference on standard models. These startups test themselves on a wider array of models at varying levels of complexity and benchmark against various chipmakers, including [NVIDIA](#), for deep neural nets, competing primarily on energy efficiency for comparable performance. Startups also focus on specific use cases, including speech detection, object detection, and facial recognition, creating competitive advantages in specific areas. In Q2 2022, edge inference startups [Quadric](#) and [Kneron](#) raised Series B rounds of \$31.0 million and \$48.0 million, respectively, demonstrating early-stage startups' ability to continue hitting milestones in a challenging commercial environment.

Inference for generative AI and 5G algorithms yields forward-looking opportunities for startups. Skyrocketing costs for large language model inference, including ChatGPT, encourage AI innovators to seek innovation in compute efficiency. Generative AI leader [Jasper](#) struck a partnership with [Cerebras Systems](#) for AI computing to diversify the company's relationships away from [OpenAI](#), [Microsoft](#), and [NVIDIA](#). In Q1 2023, AI computing startup [SambaNova Systems](#) launched a software suite of generative AI algorithms that optimize for the company's hardware. The company is offering 10 startups \$1 million in compute credits to conduct generative AI inference with the company's semiconductors. Inference-focused startup D-Matrix also sees an opportunity to benefit from generative AI by bringing inference closer to users in local data centers and workstations with AI inference within underutilized static random-access memory (SRAM). We believe that the inference market remains a Wild West for new architectures and use cases.

In 5G, startups can bring edge computing to base stations, radically reducing the power costs of high-throughput AI processing for telecom signals. In Q1 2023, edge AI startup [EdgeCortix](#) partnered with SoftBank's telecom business to use its processor for research into 5G signal processing. SoftBank previously launched the AI-to-5G Lab in partnership with [NVIDIA](#). Our briefing with photonics startup [Salience Labs](#) also revealed 5G as a leading opportunity for silicon photonics. The 5G chip market is rapidly evolving and is pushing the wide-bandgap market to \$10 billion in the medium term.⁹ Chips enhancing radio-frequency communications with signal processing can benefit from this megatrend.

Automotive semiconductors: Automotive and edge computing demands are driving more commercial agreements for inference-focused chips than for cloud training chips. In Q2 2022, edge

⁹: "Forecast Analysis: Wide-Bandgap Semiconductor Devices, Worldwide," Gartner, August 16, 2022.



AI & ML SEMICONDUCTORS

AI chip startup [Hailo AI](#) announced a partnership with leading automotive chipmaker [Renesas](#) for self-driving applications. [Renesas](#) is on pace to generate \$3.0 billion in automotive revenue in 2022 and was only the third-largest automotive chipmaker in the world in 2021, demonstrating the scale of this market.^{10, 11} The company plans to package [Hailo AI](#) neural network processors with its advanced driver assistance system chips to facilitate autonomous driving features. Startups focusing on datacenter training remain reliant on contracts with research laboratories for specific tasks.

In Q3, edge AI chip startup [SiMa.ai](#) raised a \$67.8 million Series B1 extension in part to pursue the automotive market in a strategic pivot. The company attributed the decision to the size of the automotive opportunity in computer vision. As part of the raise, the company hired veteran automotive executive Harry Kroger as president of automotive, demonstrating the urgency of the opportunity and investors' belief in the pivot.

Silicon photonics: The technology converts binary values into photons, which can more rapidly transfer data than digital bits and with less heat loss. Photonics can continue to improve AI training efficiency in the long run, thus creating the opportunity for large standalone chipmakers to be created in the emerging space. Photonics startups are achieving superior results on training speed for leading AI model types. Startup [Lightmatter](#) is optimizing its hardware for leading neural network types including ResNet for computer vision, BERT for NLP, and [Meta](#) AI's deep learning recommendation model for digital personalization. The company claims to reduce training times by 10x compared with [NVIDIA](#)'s A100, though these test results have not been verified by a

standards body.¹² We believe that this research focus demonstrates that photonics will be used for AI training in the future, as opposed to more common current use cases including interconnect and switches.

In Q1 2022, the market downturn affected silicon ASICs but did not hinder emerging photonics companies. AI chip design startups raised only \$470.7 million during the quarter after raising \$7.4 billion in 2021. Of these deals, the two largest were for silicon photonics startups. [Luminous](#) raised the largest deal we have tracked in AI-specific photonics startups, with a \$105.0 million Series A from a syndicate of investors including Bill Gates to pursue commercialization. The company has [five active patents](#), including for parallel photonic processing, based on prior research by the company's CTO at Princeton University. The company is focusing on interconnect applications in the short term with a long-term view to enable photonic processing.

Risks & considerations

The slowing of Moore's Law threatens deep learning progress: Moore's Law observes an empirical and finite relationship in which transitory density in an integrated circuit doubles every two years, but the time this actually takes has extended. The costs of continuing Moore's Law are justified only because it enables continuing chip improvements, such as transistor efficiency, transistor speed, and the ability to include more specialized circuits in the same chip. The past decade has seen improvements in transistor density slowing. Today, leading chips contain billions of transistors, but they have 15x fewer transistors than they would have if Moore's Law had

10: "3Q 2022 Presentation," Renesas Electronics Corporation, October 26, 2022.

11: "Worldwide Automotive Semiconductor Market Shares, 2021: Established Vendors Continue Dominance," IDC, Nina Turner, July 27, 2022.

12: "It's Going To Run a Lot of the Backbone of the Internet': Lightmatter CEO Nick Harris on Its Energy-Efficient AI Chips," Morning Brew, Hayden Field, July 23, 2021.



AI & ML SEMICONDUCTORS

continued unabated. For this reason, there are constraints on how much innovation is possible in raw processing power.

Chip development is expensive: The cost to design a leading semiconductor at a 7-nanometer node can reach \$300.0 million, including software, physical design, testing, and tape-out. Chip companies are accustomed to regularly releasing new chips, resulting in an astronomical burn rate for startups.

Data center startups not reaching scale: Startups focusing on datacenter training remain reliant on contracts with research laboratories for specific tasks. Multiple advanced datacenter startups list Argonne National Laboratory as a leading client, despite the federal entity's low budget for procurement each year relative to hyperscalers. We believe these research-based contracts continue to take the place of significant revenue-generating contracts with hyperscalers, thereby limiting revenue growth for startups. Pilot projects completed with hyperscalers in 2021 have not evolved into publicly disclosed commercial agreements, thus leaving startups behind on their commercial milestones for future fundraises, while the datacenter AI market decelerates from its high-growth period.

Semiconductor startups may not scale without a commercially viable architecture:

Some advanced semiconductor startups lack the ability to manufacture a cost-competitive chip due to a lack of commercially viable architecture. We believe chip designs with the potential to commercialize and compete with products on the market today must feature the following characteristics:

- Pricing of under \$1.00 per giga floating point operations per second (gigaFLOPS)
- Sufficient funding to bring a chip to production
- Interconnected architectures to optimize data flow between processing and memory
- Compatibility for training and inference software, because some chips require separate software for both processes, which has knock-on effects for software design
- A fabrication node at 5 nanometers or below
- Integrated developer kits
- Validation from cloud hosts



Appendix



APPENDIX

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

Company	Category	VC raised to date (\$M)	Post-money valuation (\$M)	Lead investor(s)
Waymo	Autonomous vehicle design	\$5,500.0	\$30,750.0	Magna International, Autonation, Silver Lake
Databricks	Database management	\$3,497.4	\$38,000.0	Counterpoint Global
Anduril	Drones	\$2,315.1	\$8,480.0	Valor Equity Partners
Nuro	Industrial robots	\$2,132.0	\$8,600.0	Tiger Global Management
Horizon Robotics	Processor design	\$1,900.0	N/A	Chery Automobile
Indigo Agriculture	Crop maximization	\$1,701.6	\$3,950.0	Empede Capital Partners, 369 Growth Partners, Timon Capital, Molini Besozzi Marzoli
ContentSquare	Ad targeting	\$1,412.6	\$5,686.4	Sixth Street Partners
SambaNova Systems	Processor design	\$1,136.6	\$5,000.0	SoftBank Investment Advisers
DataRobot	AlaaS	\$1,048.1	\$6,300.0	Altimeter Capital Management, Tiger Global Management
OpenAI	AlaaS	\$1,010.1	N/A	Microsoft

Source: PitchBook | Geography: Global | *As of December 31, 2022



APPENDIX

Top strategic acquirers of AI & ML companies since 2017*

Name	Deal count
ServiceNow	5
Oracle	5
Accenture	5
Apple	5
International Business Machines	4
Meta Platforms	4
Intel	3

Source: PitchBook | Geography: Global | *As of December 31, 2022

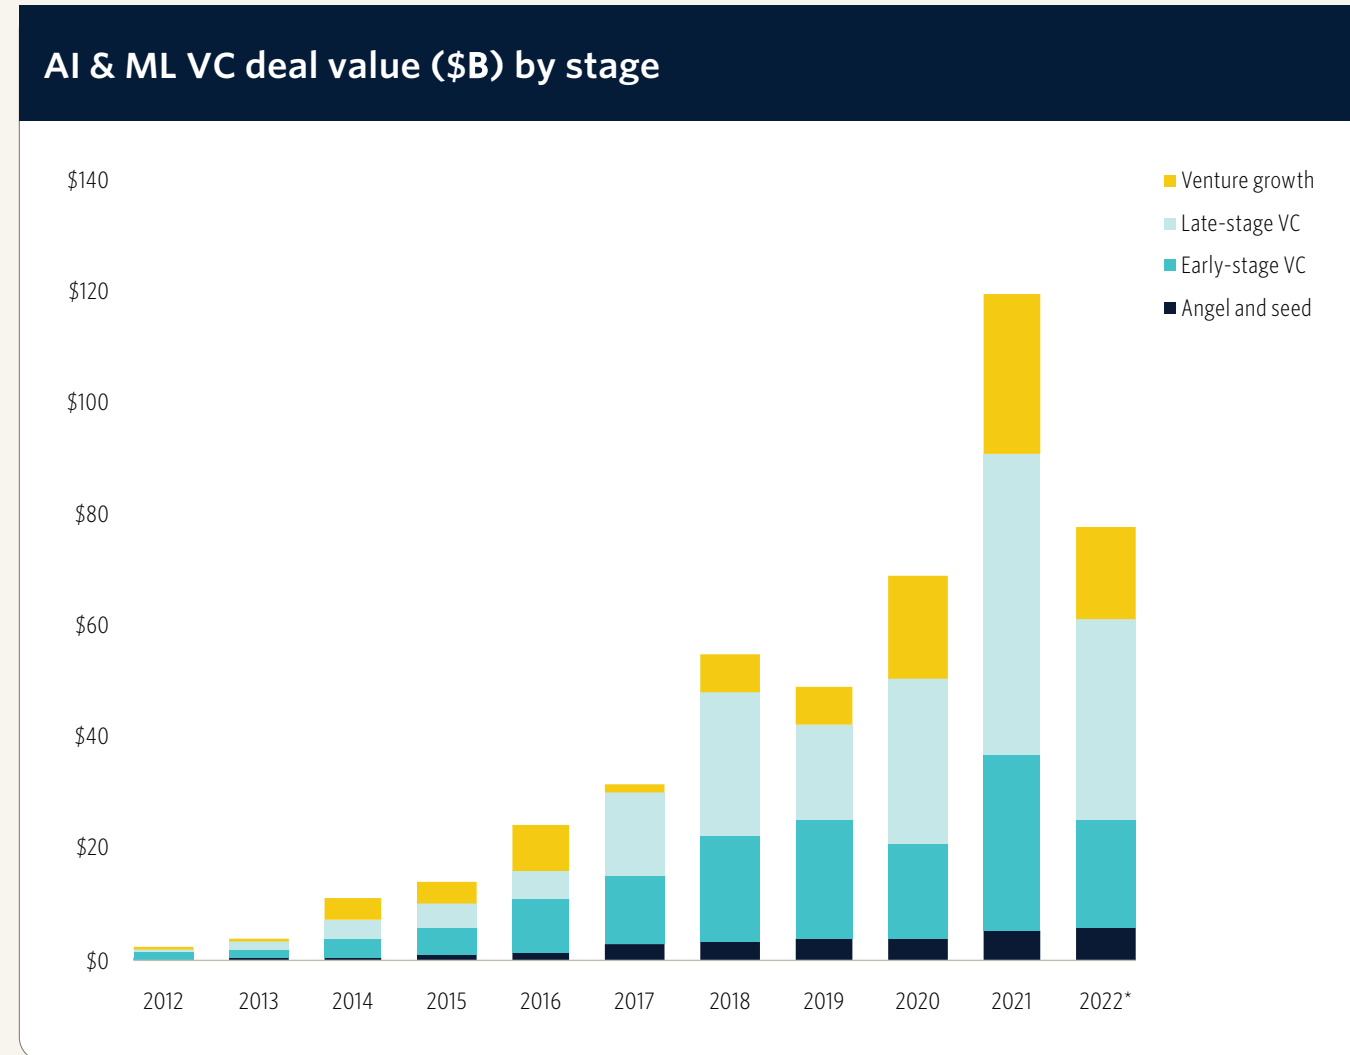
Top VC investors in AI & ML companies since 2017*

Name	Deal count
Alumni Ventures	266
Sequoia Capital China	245
SOSV	202
500 Global	176
Enterprise Ireland	174
Right Side Capital Management	169
Khosla Ventures	168
Accel	168
Sequoia Capital	167
Tiger Global Management	167


Source: PitchBook | Geography: Global | *As of December 31, 2022



APPENDIX




Source: PitchBook | Geography: Global | *As of December 31, 2022




Source: PitchBook | Geography: Global | *As of December 31, 2022



APPENDIX



Source: PitchBook | Geography: Global | *As of December 31, 2022



Source: PitchBook | Geography: Global | *As of December 31, 2022

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
SaaS

Ali Javaheri
ali.javaheri@pitchbook.com
Emerging Spaces

Robert Le
robert.le@pitchbook.com
Web3/DeFi
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
Fintech

Coming soon
Biotech
Pharmatech