


EMERGING TECH RESEARCH

Vertical Snapshot: Generative AI

VC trends, industry overview, and market landscape



2023



Contents

Executive summary	3
Introduction	4
Technology and processes	5
Tech stack	7
Timeline	8
Industry drivers	12
Market size	14
Outlook	15
Considerations	17
Taxonomy	18
Market map	23
VC activity	25
Key players	29
Key partnerships	33
Key acquisitions / deals	34
Glossary	37

Institutional Research Group

Analysis



Ali Javaheri Associate Analyst, Emerging Technology
ali.javaheri@pitchbook.com
pbinstitutionalresearch@pitchbook.com



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

Data

Matthew Nacionales, Data Analyst

Publishing

Report designed by **Drew Sanders**

Published on March 22, 2023



Executive summary

- **VC investments soar in natural language interfaces, 2D media, AI Core, and biotech startups.** Between 2018 and 2022, natural language interfaces and 2D-media generation captured 29.3% and 18.8% of VC-backed deals, respectively, and 23.6% and 11.4% of total dollars invested, respectively, in the generative AI space, making them the leading applied segments of the technology. Within the natural language interfaces segment, chatbots, voicebots, and personal assistants captured \$544.9 million in 2022—or 59.6% of all dollars invested in the segment. For the 2D media segment, avatars and video generation and editing accounted for 37.7% and 40.8%, respectively, of all VC dollars invested in the segment in 2022. AI Core, which includes foundation model developers and infrastructure for model development, raised a whopping \$5 billion from VCs between 2018 and 2022. Remarkably, even biotech startups utilizing generative techniques have been hot destinations of VC capital—with \$1.6 billion invested during the same period.
- **Lower training costs and advances in research propel in-house teams and startups.** The declining cost of foundation model training and usage, along with advances in artificial intelligence & machine learning (AI & ML) research, is driving the industry forward, making it possible for both in-house research teams and research-oriented startups to thrive. With hardware optimization and customized software, competitive foundation models can now be trained at significantly lower costs. As more AI researchers move from academia to the private sector, companies will compete for talent, and startups that simultaneously dedicate resources to model and product development will be more competitive.
- **Hyperscalers investing heavily in startups.** Cloud providers recognize that leading innovation in generative AI algorithms comes from startups; these providers are both making significant acquisitions and striking generous partnerships to align themselves with the future of creative software. Microsoft's \$10.0 billion financing commitment for OpenAI and Alphabet's VC mega-deals for Anthropic and Cohere stand out for their scale, yet Meta, Spotify, and Apple have made significant acquisitions in generative audio as well. These investments come after a relatively quiet period for hyperscaler acquisition activity in AI, demonstrating that tech giants now require urgency to prevent disruption.
- **New business opportunities from generative foundation models will rely on user experience.** Generative foundation models are still in their infancy, but this presents an exciting chance for basic innovation to produce better user experiences. We have already seen commercially successful products based on generative transformer and score-based diffusion models, so a solid precedent exists for future research and development to yield results within just 18 to 24 months. And as model capabilities continue to advance, new business models and applications will likely emerge, including exciting experiments with video, voice, scientific publications, and databases. However, it is crucial to ensure that foundation models have guardrails in place to guarantee trustworthy outcomes. For startups looking to take on industry incumbents, the keys to success will be user experience, customization, and access to proprietary data.
- **The generative AI market is expected to grow, with a multitude of enterprise applications.** Generative-relevant use cases already present a significant enterprise opportunity, estimated to reach \$42.6 billion in 2023, with natural language interfaces offering the largest market due to customer service and sales automation use cases. We expect the market at a 32.0% CAGR to reach \$98.1 billion by 2026, even without accounting for the potential of generative AI to expand the total addressable market of AI software to consumers and new user personas in the enterprise.



Introduction

Imagine a world where machines can create art, write articles, design products, and even hold conversations that are virtually indistinguishable from those with a human being. This futuristic vision is fast becoming a reality thanks to the rapid advances in the field of generative AI. Tools such as ChatGPT and Stable Diffusion have captured the public's imagination and showcased the potential of this technology to revolutionize numerous industries. But what is generative AI, exactly? In this report, we define generative AI to encompass all users and enablers of content produced by neural networks. Its applications are incredibly broad in scope. From creating engaging content and automating business processes to developing software, improving customer service, and even advancing biotechnology, generative AI is poised to significantly augment the way we live and work. The way we deploy AI will also have to change given the advanced models and new interfaces underlying this trend.

Investors have taken note, with VC activity growing exponentially since the release of OpenAI's GPT in 2018. While VCs invested \$408.0 million that year, the numbers have since skyrocketed, with \$4.8 billion and \$4.5 billion invested in 2021 and 2022, respectively, despite a recently depressed market. Angel and seed deals have also grown, with 107 deals and \$358.3 million invested in 2022 compared with just 41 and \$102.8 million in 2018. The surge in interest from early-stage VC companies has been staggering, with a total of \$2.2 billion raised in 2022 alone. Big Tech companies are partnering with research startups to drive innovation, and with the cost of foundation models and training coming down, other startups are finding new opportunities to build novel applications on top of foundation models.

MidJourney-generated Image



Prompt used: Isfahan, cubism



Technology and processes

Generative model architectures¹

Transformer: Transformers are a type of AI architecture that can process and generate sequences of data, such as text or speech. They use attention mechanisms to allow the model to focus on different parts of the input sequence at each step, and they can generate output sequences that are conditioned on the input sequence. They have been used to create state-of-the-art language models and to generate realistic-sounding speech. Perhaps the most prominent example of an application built upon a transformer model is OpenAI's ChatGPT.

Diffusion: Diffusion models are a type of generative model that can create high-quality images by gradually adding and removing noise from an image in repeated steps.² They are based on the idea of systematically and slowly destroying structure in a data distribution through an iterative forward diffusion process and then learning a reverse diffusion process that restores structure in data. Examples of diffusion models include OpenAI's Dall-E and Stability.ai's Stable Diffusion. They are one of the most prominent architectures in generative AI today.

Variational autoencoders (VAEs): These models can learn to create new data by compressing existing data into a lower-dimensional space and then generating new data from that space.³ They consist of two parts: an encoder that compresses the data and a decoder that generates new data from the compressed representation. The encoder is trained to produce compressed data that can be used to generate new data that is similar to the original data. The decoder is trained to generate new data that is similar to the compressed data, but also has some randomness added to it to make it more varied.

Generative adversarial networks (GANs): GANs produce new data that resembles a given dataset. These models consist of two parts: a generator that creates new data and a discriminator that evaluates whether the data is real or fake. The two parts are trained together in a game-like fashion, with the generator trying to create data that can fool the discriminator into thinking it is real, and the discriminator trying to correctly discern real data from fake data. Over time, the generator learns to create data that is more and more like the real data, until it is difficult to tell the difference. They have historically been used to produce deepfake images.

Of these models, transformer and diffusion are the most prominent in generative AI today.

1: "Generative AI: All You Need to Know," MURF AI, January 6, 2023.

2: "Generative AI I: Variational Autoencoders" Synthesis Blog, February 7, 2023.

3: "Introduction to Diffusion Models for Machine Learning" AssemblyAI, Ryan O'Connor, May 12, 2022.



TECHNOLOGY AND PROCESSES

Foundation models

Foundation models are powerful models that are trained on large datasets with minimal data labeling, allowing them to adapt to a wide range of downstream tasks. These models leverage advances in deep learning, particularly transfer learning and self-supervised learning. Large language models (LLMs) are one example of foundation models along with image generation, audio, and protein folding models.⁴ Foundation models make connections between diverse datasets and can respond to user inputs. Text-based models can respond to natural language instructions that include references to specific variables, such as the sentiment of text, genre of music, or style of image. Some foundation models are even multimodal, meaning they are trained on a variety of data types.

For further analysis of foundation model architectures, please refer to our [recap of Fortune's Q4 2022 Brainstorm AI Conference](#).

4: "On the Opportunities and Risks of Foundation Models," Stanford University Center for Research on Foundation Models, August 16, 2021.

Emerging approaches to foundation model architecture

Centralized foundation models

Trained by research labs for distribution by vendors

Can be tuned for specific applications

Based on general purpose dataset

Leverage state-of-the-art techniques to maximize accuracy

Customized multi-modal foundation models

Trained by AI vendors in partnership with customers

Designed to be application-specific for multiple customers

Combine general purpose datasets and niche datasets

Use mixed techniques to balance accuracy and efficiency

Customer-specific foundation models

Trained by vendors or customers

Designed to be application-specific

Based on niche datasets

Use simplest techniques to maximize efficiency

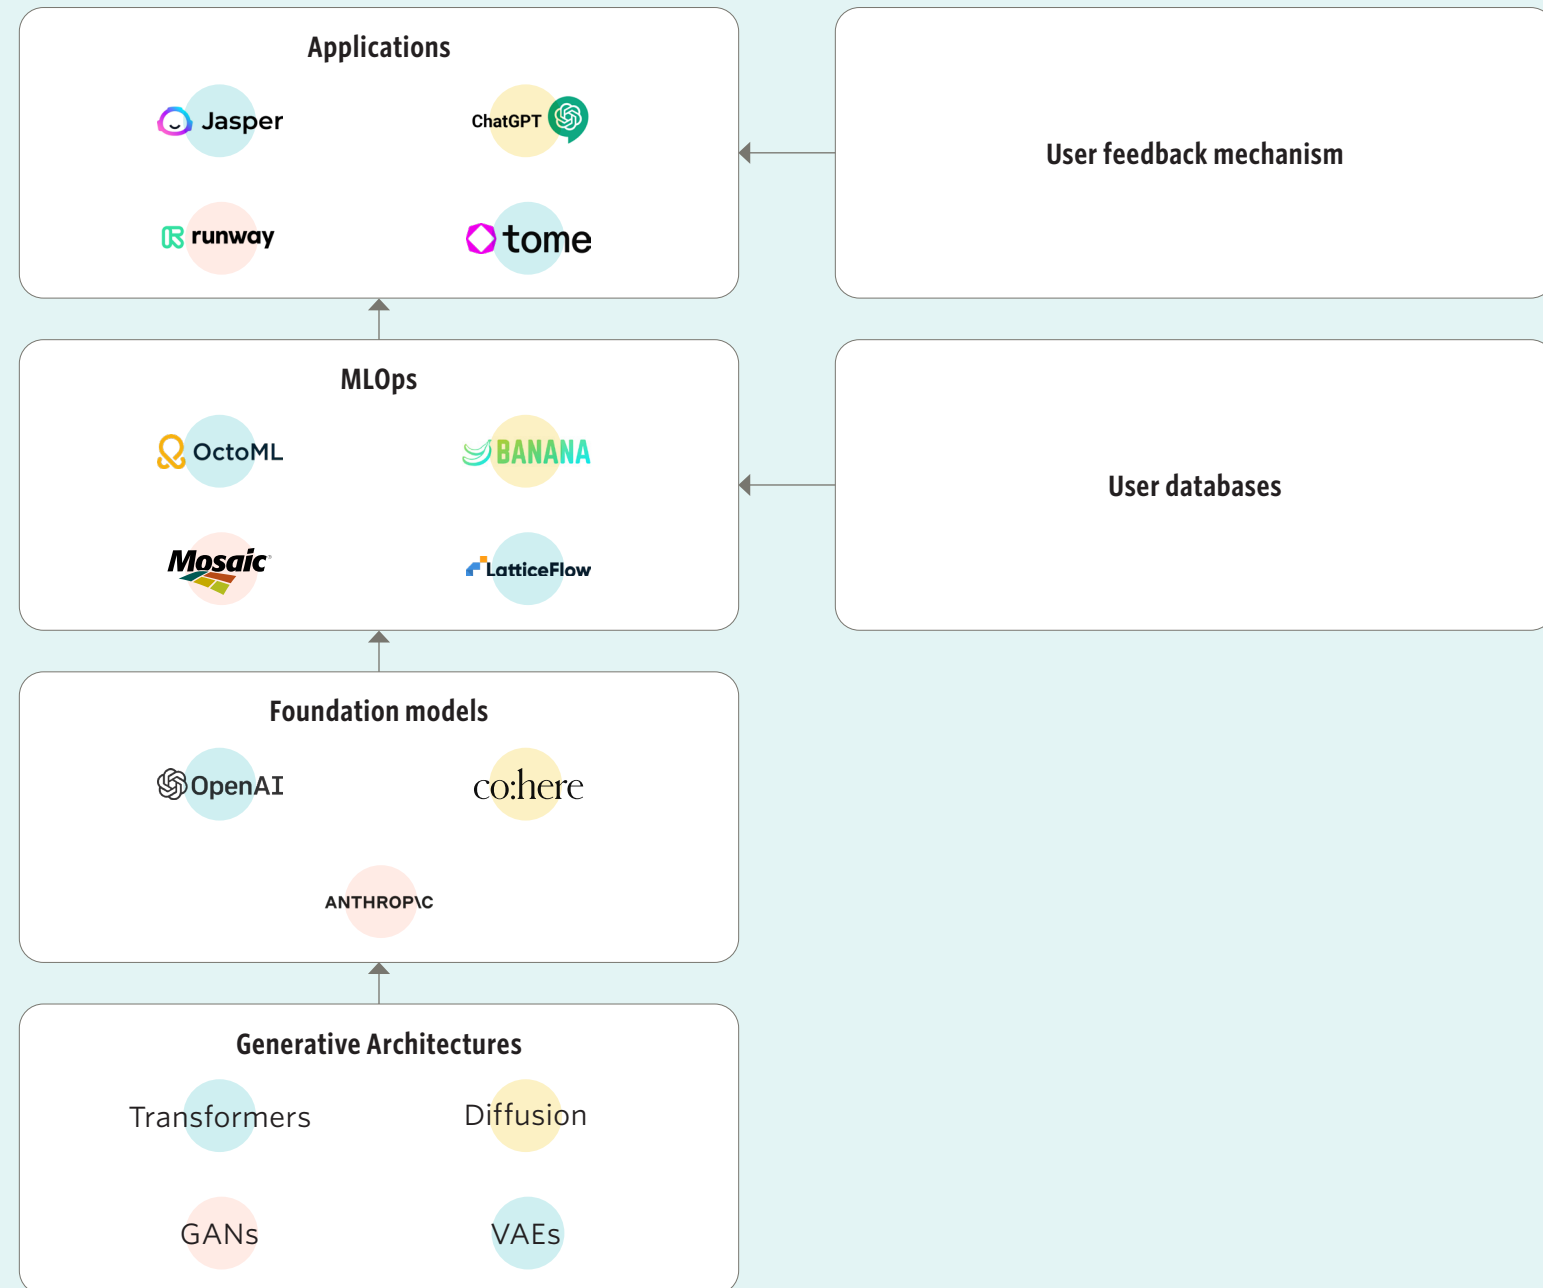
Source: PitchBook



TECHNOLOGY AND PROCESSES

Pre-trained foundation models can be “fine-tuned” for specific use cases by using a domain-specific dataset to supplement the broader dataset and any manual rule-setting of model outputs. Fine-tuning enables the development of new applications on top of foundation models, thereby reducing the cost of creating AI-enabled products. Models can be connected to application databases to learn from user actions and learn connections between business logic and data types, making large language models more accurate than looking through public resources for information. Once applications are deployed, users can offer feedback on model outputs for continued refinement of model guardrails.

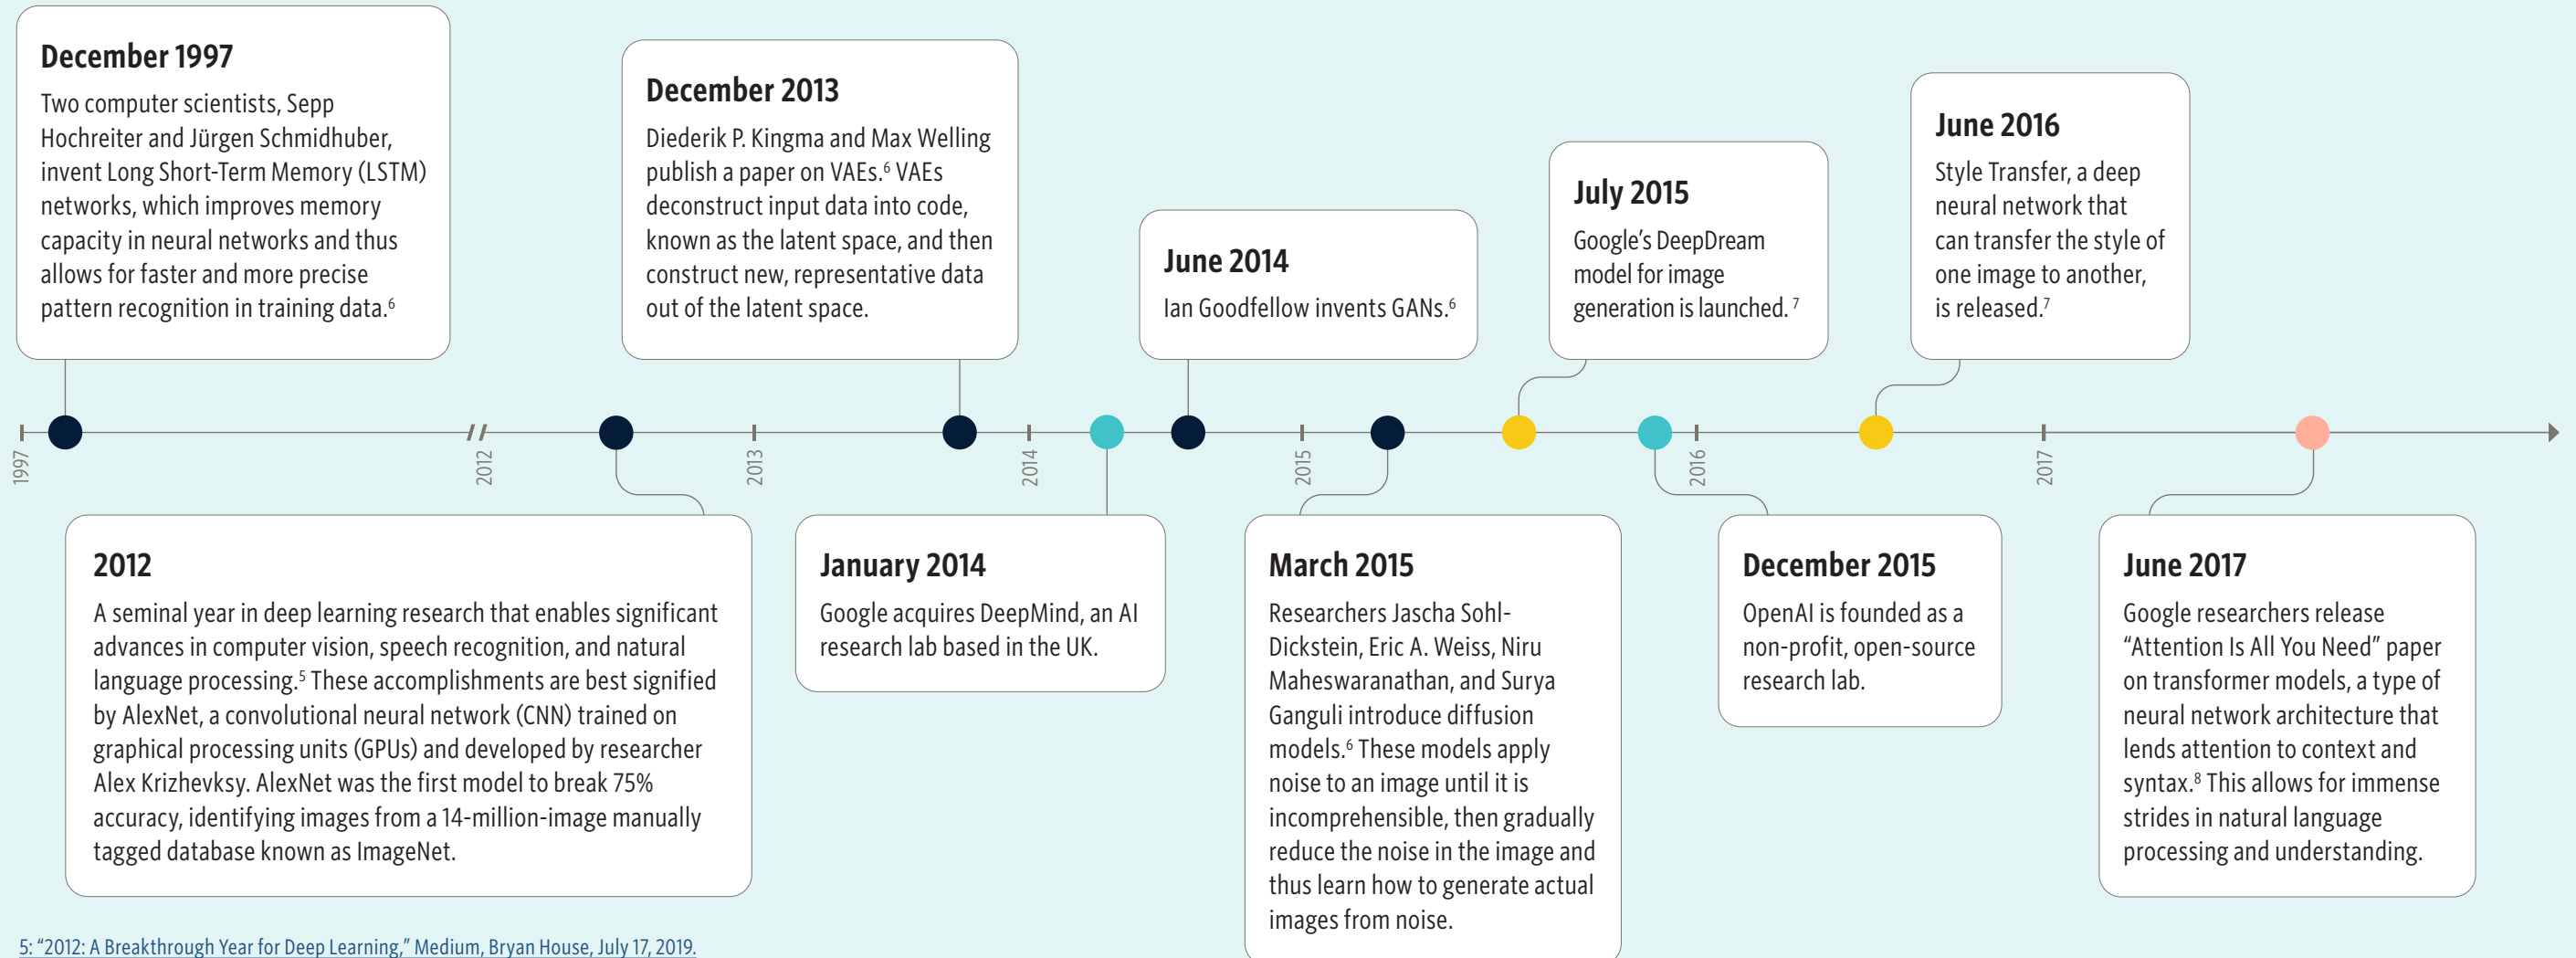
The generative AI tech stack





Generative AI timeline

● General model research ● Image / video models ● Text models ● Business events



5: "2012: A Breakthrough Year for Deep Learning," Medium, Bryan House, July 17, 2019.

6: "The Generative AI Revolution Has Begun—How Did We Get Here?," Ars Technica, Haomiao Huang, January 30, 2023.

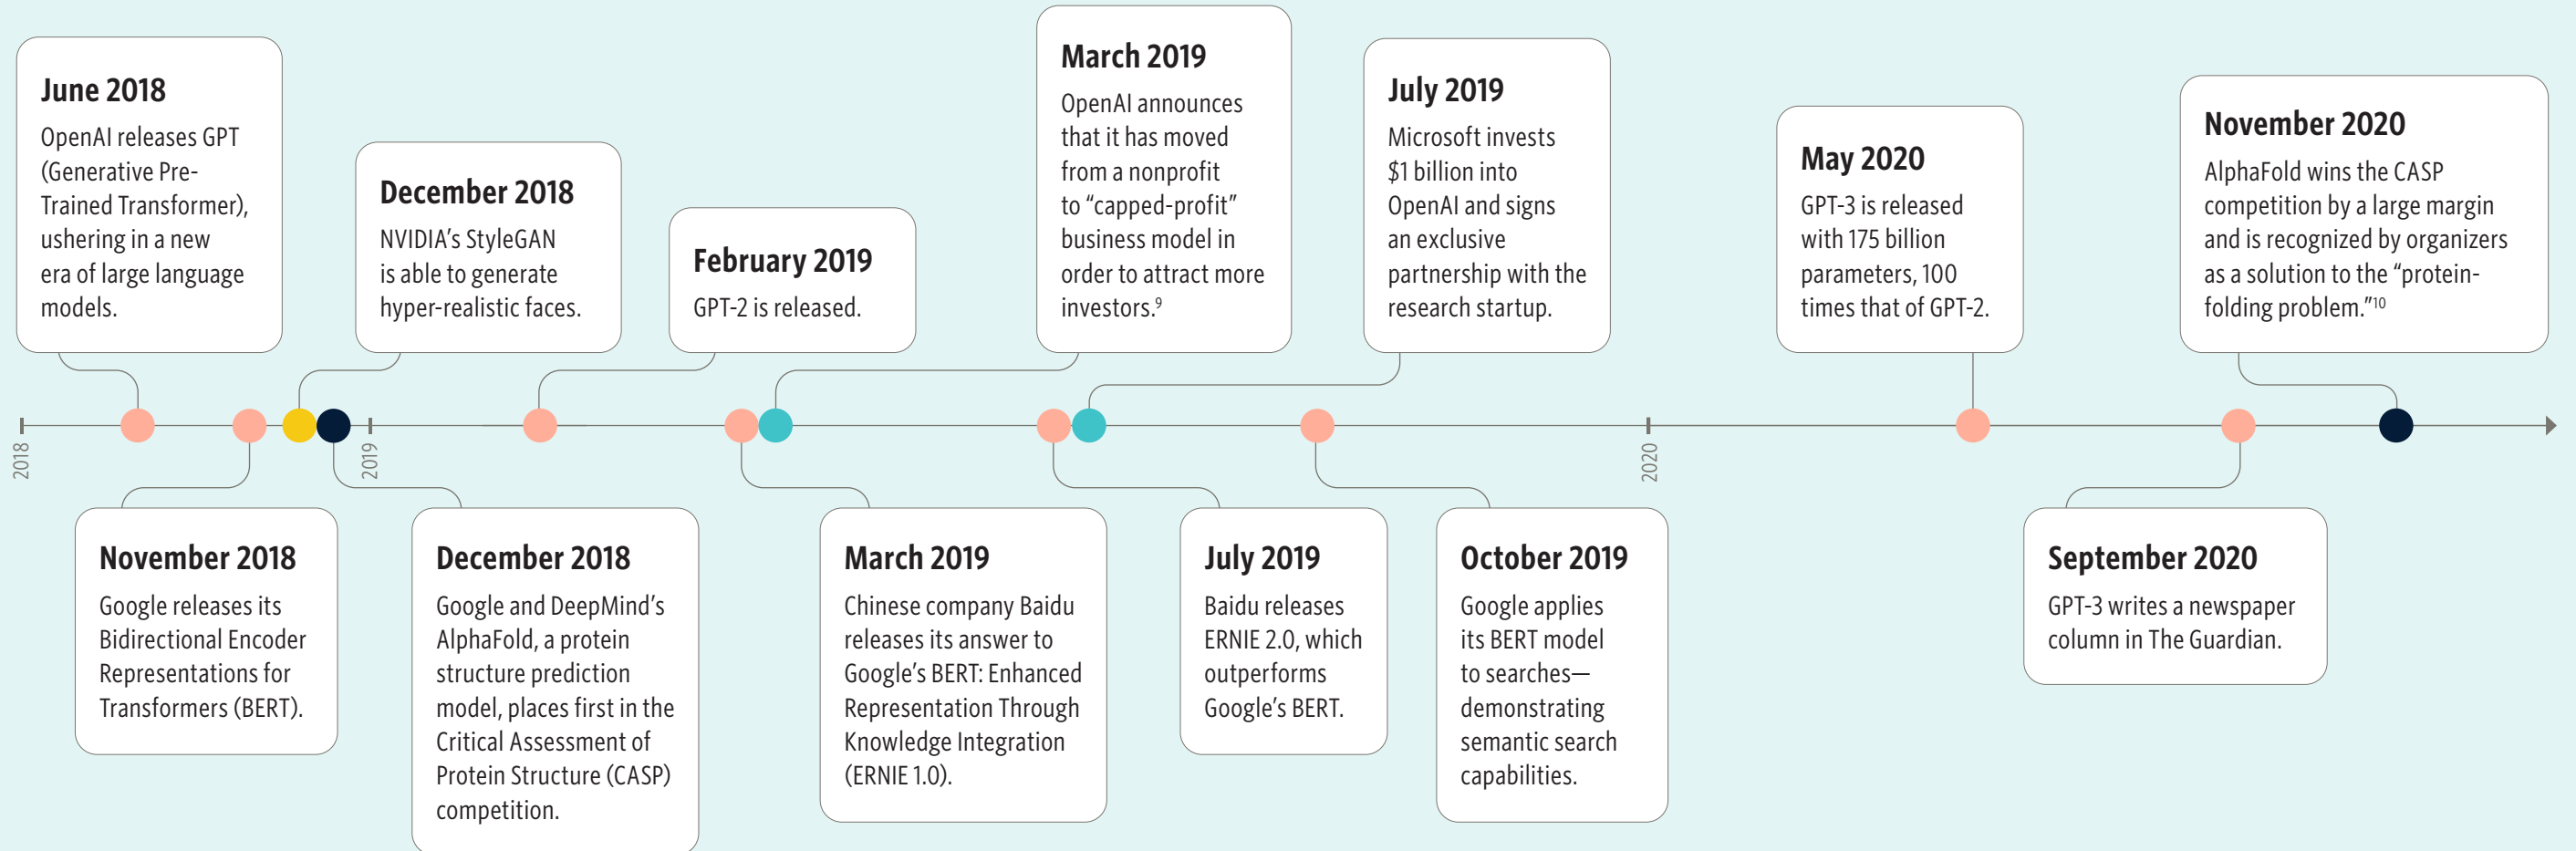
7: "Timeline of Text-To-Image Machine Learning Models," Fabian Mosele, n.d., accessed March 9, 2023.

8: "Timeline of AI and Language Models," Life Architect, Alan D. Thompson, n.d., accessed March 9, 2023.



GENERATIVE AI TIMELINE

● General model research ● Image / video models ● Text models ● Business events



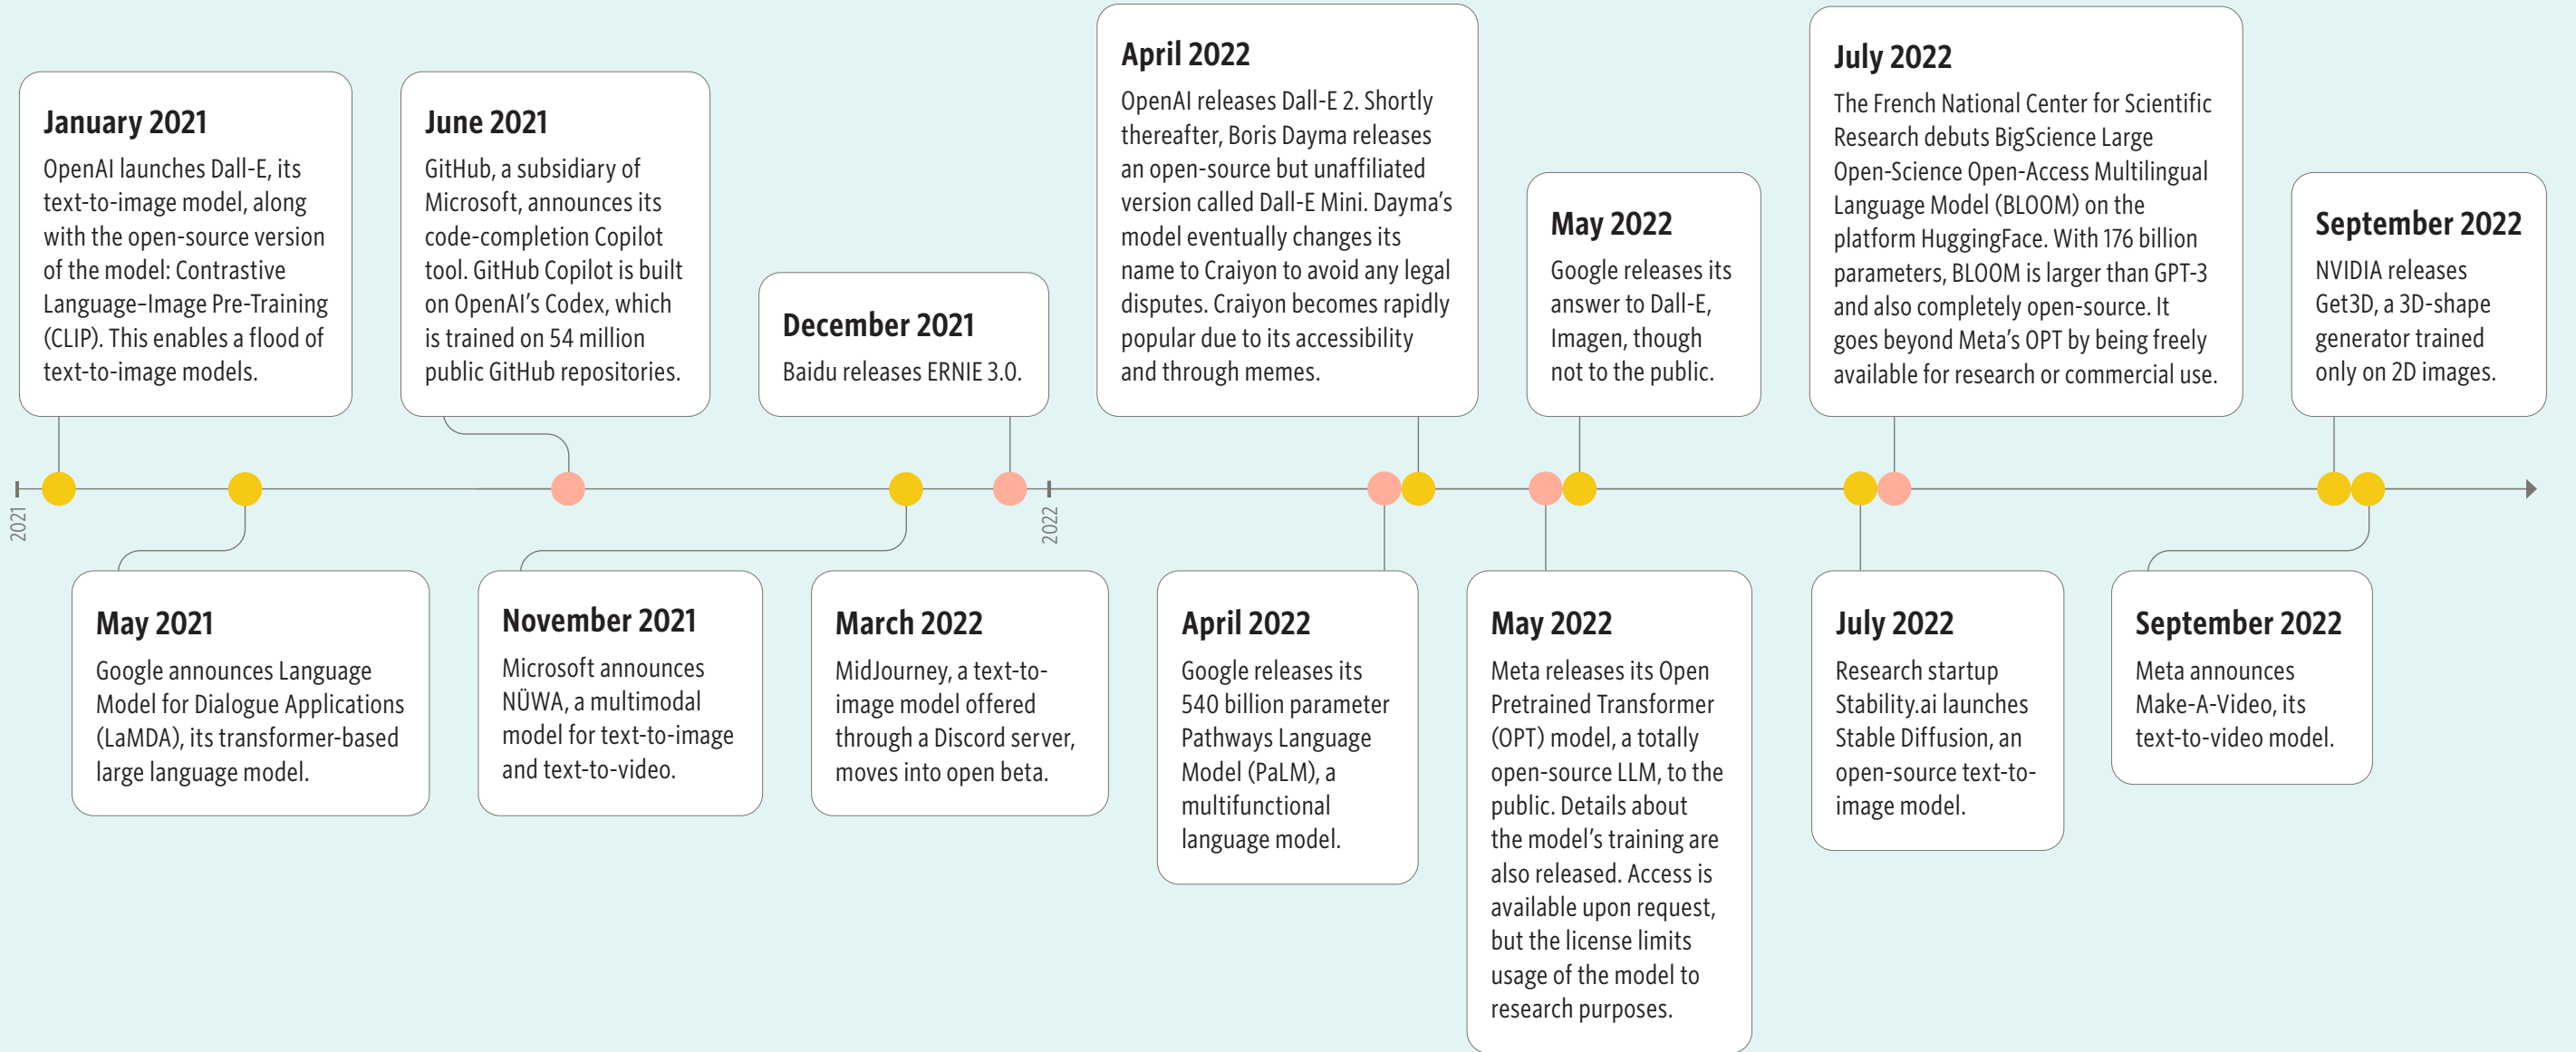
9: "OpenAI Shifts From Nonprofit to 'Capped-Profit' to Attract Capital," TechCrunch, Devin Coldewey, March 11, 2019.

10: "DeepMind's Protein-Folding AI Has Solved a 50-year-old Grand Challenge of Biology," MIT Technology Review, Will Douglas Heaven, November 30, 2020.



GENERATIVE AI TIMELINE

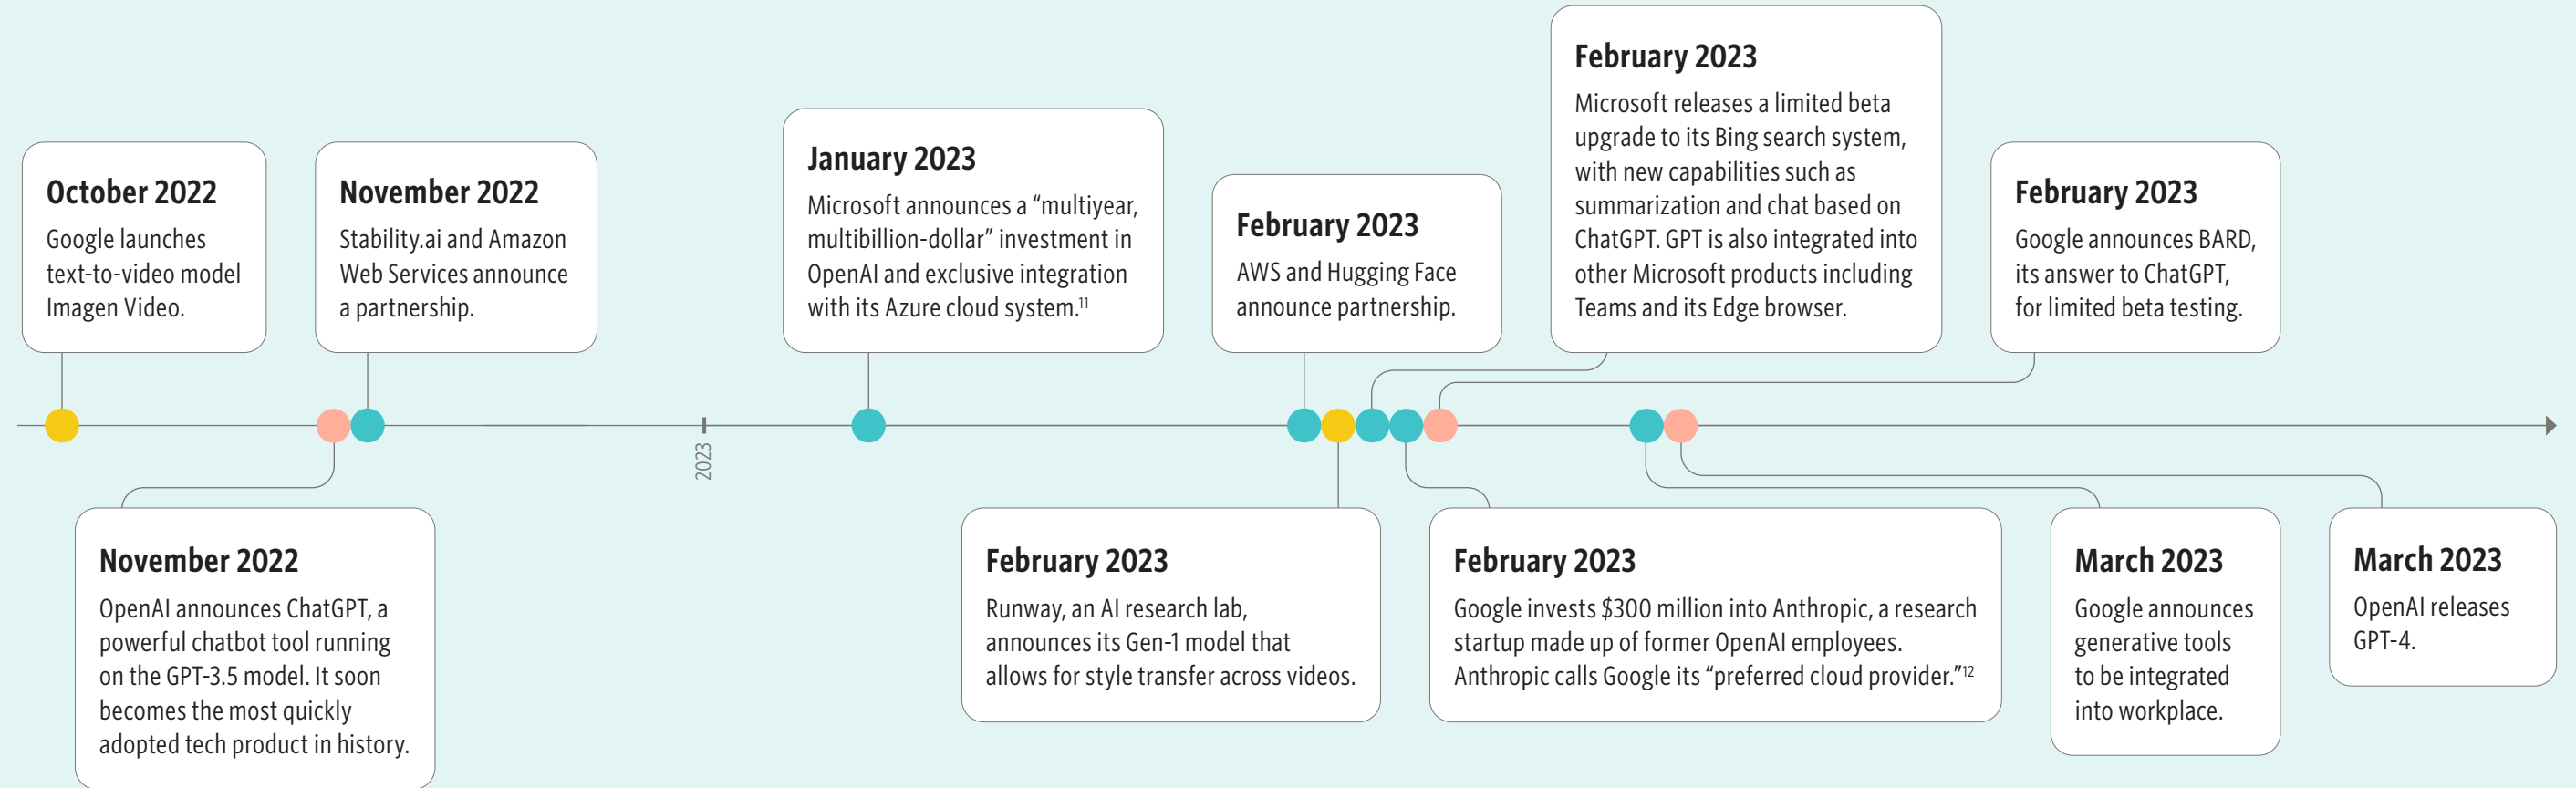
● General model research ● Image / video models ● Text models ● Business events





GENERATIVE AI TIMELINE

● General model research ● Image / video models ● Text models ● Business events



11: “Microsoft Announces New Multibillion-Dollar Investment in ChatGPT-Maker OpenAI,” CNBC, Ashley Capoot, January 23, 2023.

12: “Google Invested \$300 Million in AI Firm Founded by Former OpenAI Researchers,” The Verge, James Vincent, February 3, 2023.



Industry drivers

Costs of foundation model training coming down: GPT-3 was 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 \$600,000.¹³ Hardware optimization can produce further cost savings, with competitive foundation models able to be trained for 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.

Costs of foundation model use remain accessible to startups: OpenAI competitively prices its models for startups to launch applications via an application programming interface (API). The cost of a typical response from OpenAI's ChatGPT costs around \$0.002, thus enabling startups to build products around generative AI. OpenAI also offers a low-cost version of its model for use cases requiring less accuracy and a reduced API that has lowered the cost of entry. In image generation, Midjourney offers pro plans that run 12 concurrent high-speed inferences for \$48 per month, making it accessible even to large enterprises.


Chief information officers (CIOs) pushing for AI transition across departments: 60% of CIOs plan for AI to gain widespread use across departments by 2025.¹⁴ In 2022, only IT, supply chain, and finance departments saw widespread adoption in more than 40% of organizations.¹⁵ The facilitation of investment by other large functions including sales, product development, and human resources can encourage generative AI adoption.

13: This estimate is a self-disclosure from Stability AI founder Emad Mostaque.

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

15: Ibid.

Affiliation of research teams building notable AI systems



Source: [Sevilla, et al.](#) | Geography: Global | *As of March 9, 2022



INDUSTRY DRIVERS

Supercharged productivity can solve labor shortages and address cost-cutting measures by businesses facing revenue headwinds: Generative tools that serve as copilots—chatbots, coding assistants, and knowledge organizers—can automate repetitive and time-consuming tasks, freeing employees to focus on higher-level work and leading to increased output per worker. As businesses address worker shortages and pressures to reduce cost, many may turn to generative AI as a solution. Companies such as Citadel are reportedly negotiating enterprise-wide ChatGPT licenses because of the productivity increase,¹⁶ and studies have confirmed the increased worker productivity across multiple domains.¹⁷

In-house research teams drive capabilities: AI researchers have steadily moved from academia to industry, and we expect this trend to continue, given the growing opportunities within the private sector. With that said, companies will compete for talent, and those product-specific startups with dedicated AI & ML researchers will be more competitive. Companies such as Runway and Synthesia already dedicate resources to model and product development simultaneously and symbiotically.

Viral consumer adoption of new products: ChatGPT reached 100 million monthly active users faster than any technology productivity app at two months, surpassing TikTok’s nine months to achieve the same feat, according to UBS analysis.¹⁸ AI slide deck software startup Tome claims to have achieved the fastest time to 1 million users for a productivity app.¹⁹ Similarly, the success of Stable Diffusion and Midjourney primed consumers to learn about new AI freeware and refer friends through social media and word-of-mouth.


16: [“Citadel Negotiating Enterprise-Wide ChatGPT License, Griffin Says,” Bloomberg, Katherine Doherty and Felipe Marques, March 7, 2023.](#)

17: [“Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence,” MIT Economics, Shakked Noy and Whitney Zhang, March 2, 2023.](#)

18: [“Let’s Chat About ChatGPT,” UBS Investment Research, February 23, 2023.](#)

19: [“Tome Is Fastest Productivity Tool to Reach 1M Users,” GlobeNewswire, February 22, 2023.](#)

Dall-E-2-generated image




Prompt used: A curious robot reading an interesting market research report about robots, in the style of Basquiat



Market size

We estimate that generative AI addresses an existing enterprise opportunity and will reach \$42.6 billion in 2023. This includes budgets that enterprises are already dedicating to third-party AI software, including AI Core and applications. Among our generative AI segments, natural language interfaces offers the largest market due to the presence of customer service and sales automation use cases. Vertical applications and AI Core offer about the same revenue opportunity, given that not all vertical applications can be addressed with generative AI. We include AI-enabled business automation software along with other minor categories, including biotechnology, human resources, and fraud-detection software. Visual content generation is a small market relative to natural language processing for enterprises, and we do not assume it will replace conventional graphic design software. Even without accounting for the potential of generative AI to expand the total addressable market of AI software, this current market size is on pace to reach \$98.1 billion by 2026, growing at a 32.0% CAGR.

AI software market size (\$B) estimates in generative-relevant categories*



Source: PitchBook Emerging Tech Research | Geography: Global | *As of August 26, 2022



Outlook

Given the relative immaturity of generative foundation models, basic innovation can yield fundamentally improved user experiences. Generative transformer models and score-based diffusion models emerged in 2020 and 2019, respectively, and have since yielded commercially viable products. Based on this precedent and the recent rollout of developer-friendly APIs for leading large language models, we believe that basic research & development (R&D) can be commercialized within 18 to 24 months. Continued innovation in model capabilities can create new business models as current models specialize in certain functions such as text-to-image and question answering.

New models from research labs may make new product types and applications possible in diverse creative and knowledge-based domains. Experiments with video, voice, scientific publications, and databases could make their way into foundational language models and enable new forms of learning. AI currently struggles to turn language-based instructions into polished final products in visual and knowledge-based domains. Additional samples from across visual domains including video and completed professional documents can help produce more useful outputs from few-shot prompting. Incremental data types may include practical workflows that can help AI learn how humans reason. The current distinction between large language models and text-to-image models may collapse as research labs incorporate diverse training datasets in similar transformer training runs. Products with multimodal functionality will set themselves apart from singular applications—and be better poised to compete with enterprise behemoths such as Microsoft and Adobe.

Foundation models will require guardrails to produce trustworthy outcomes. The application of reinforcement learning to GPT-3 made the ChatGPT product a partially trustworthy assistant by rewarding valid outputs through manual review. This method can benefit other models such as image generation yet remains secondary to the model training process. Techniques to fine-tune models based on end-user data will also be central to their adoption. Recent papers show that foundation models can improve in niche domains such as financial data when presented with domain-specific information.²¹ A tooling layer may emerge to guide the training, fine-tuning, and deployment of generative applications as it has for supervised machine learning. Given the differences between legacy ML analytics and generative AI interfaces, new forms of tooling will be required.

The future lies with proprietary data, fine-tuning, and user interfaces. Startups seeking to disrupt incumbents and move beyond simply providing a “wrapper” over foundation models will need to focus on user experience and customization. In this regard, vertical applications, or those for specific use cases, may stand out—already, activity in generative biotechnology is promising. Those with access to proprietary data—be they healthcare data or corporate knowledge bases—may find it easier to provide tailored solutions to clients and capture more market share. A user interface that provides easy opportunity for humans to input feedback, and robust learning mechanisms that incorporate feedback, will propel startups forward.

²¹: [“Sentiment Spin: Attacking Financial Sentiment with GPT-3,” Swiss Finance Institute Research Paper, Markus Leippold, January 30 2023.](#)



OUTLOOK

Users also want models to connect to the internet to learn from real-time data and power applications they already use. Learning from real-time data remains experimental as shown by the recent launch of Bing's chatbot, Sydney. Startups are also integrating live internet data into search applications with varying results. In 2022, Meta, NVIDIA, and OpenAI combined large language models with optimized forms of reinforcement learning for specific tasks such as the games Diplomacy and Minecraft, which enabled models to complete tasks in real-time based on imitation of user workflows. If models know about the contents of videos showing user workflows and understand internet tutorials, startups can better apply imitation learning to carry out specific tasks across domains.

Midjourney-generated image



Prompt used: Asiatic cheetah Persian miniature style



Considerations

As generative AI continues to make impressive strides, concerns have arisen among lawmakers regarding potential issues such as the production of misinformation, privacy concerns, and intellectual property protections. However, different jurisdictions have taken differing approaches toward AI regulation. China, for example, has implemented regulations around “deep synthesis technologies,” mandating watermarks, algorithm registry, and requiring user consent for data usage by companies developing these technologies.²¹

Meanwhile, the EU has classified foundation models as “high risk” under its proposed Artificial Intelligence Act, while chatbots and synthetic content generators are deemed “limited risk.”²² It remains to be seen whether companies developing applications using foundation models must abide by limited or high-risk regulations. In the UK, proposals suggest that security and transparency will underpin regulation of generative models, with a differentiated and context-specific regulatory framework that emphasizes proportionate, “pro-innovation” approaches.²³

In the United States, self-regulation and nonbinding risk-management frameworks have generally been the approach, with a particular focus on security and transparency.²⁴ In general, AI model regulations do not appear to be a primary campaign concern. However, intellectual property issues have generated media attention, with OpenAI, GitHub, and Stability.ai all being sued for allegedly infringing upon the intellectual property of developers, photographers, and artists. While the issue of fair use and copyright protection for AI-generated content is still being debated, it is unlikely that it will significantly impact the operations of platform companies, regardless of the outcome of lawsuits.²⁵ We believe that companies taking responsible approaches to data collection and governance can gain an advantage over risk-taking first movers in the long run.

²¹: [“AI 101: The Regulatory Framework,” Lewis Silkin, February 20, 2023.](#)

²²: [“Regulatory Framework Proposal on Artificial Intelligence,” European Commission, n.d., accessed March 9, 2023.](#)

²³: [“Establishing a Pro-Innovation Approach to Regulating AI,” the Secretary of State for Digital, Culture, Media and Sport, July 20, 2022.](#)

²⁴: [“Blueprint for an AI Bill of Rights,” the White House Office of Science and Technology Policy, October 2022.](#)

²⁵: [“Microsoft, GitHub, and OpenAI Ask Court to Throw Out AI Copyright Lawsuit,” The Verge, Emma Roth, January 28 2023.](#)



Generative AI taxonomy

This overview is designed to provide a high-level survey of the emerging space of generative AI, acknowledging the breadth of its applications across industries. Our [Q3 2022 Artificial Intelligence & Machine Learning Report](#) included only foundation model developers with commercially available models and generative-AI-native companies. Rapidly, companies are pivoting to support generative AI development and insert the technology at the core of their products, expanding the ecosystem. In this taxonomy, we define the infrastructure enablers and primary applications of these techniques as of publication date. The listed models in each segment are prominent examples, not exhaustive lists, and not all companies in the space will necessarily use them. To help illustrate the types of companies operating within each subsegment, representative companies are listed as an example in each category.

Stable-Diffusion-generated image



Prompt used: A frog with a top hat and monocle



GENERATIVE AI TAXONOMY

2D media		
<p>Companies in this segment produce, manipulate, or identify flat digital content—images, videos, or other media that lack spatial depth. These tools enable use cases in marketing, design, entertainment, and e-commerce.</p> <p>Models: DALL-E (OpenAI), Midjourney, Stable Diffusion (Stability.AI), Make-A-Video (Meta), Imagen Video (Google)</p>		
Avatars	Digital representations of individuals or groups of people used in personalized videos, virtual assistants, or e-commerce.	Soul Machines , Synthesia , DeepBrain AI
Content suite	Automated content creation and optimization of multiple forms of digital media for content marketing, advertising, and presentations.	Lightricks , Jasper , Tome
Image: Generation and editing	Creation, modification, and enhancement of digital images for graphic design, photography, and e-commerce.	Imagen , PhotoRoom
Product design	Assistance with the design and development of physical products such as clothing and furniture using generative design techniques.	ZMO.AI , CALA
Video: Generation and editing	Automated video editing, enhancement, and generation for video production, advertising, and entertainment.	Runway , VideoVerse , Waymark
Visual search	Automated identification and categorization of visual patterns in images and videos for e-commerce, content discovery, and analytics.	Air , Twelve Labs

3D media		
<p>Companies in this segment produce digital content represented and viewed in three dimensions. This type of media is often used in virtual reality, gaming, product design, and architecture. Generative 3D media enables immersive, interactive experiences simulating the real world.</p> <p>Models: GET3D & Magic3D (Nvidia), Multiview Compressive Coding (Meta), Point-E (OpenAI)</p>		
Avatars and nonplayer characters (NPCs)	Digital representations of individuals or groups of people or non-playable characters used in virtual worlds, gaming, and training simulations. Includes representations developed through motion capture technology.	rct AI , Inworld AI , Move.AI
Object synthesis	Automated creation and modification of 3D objects for use in video games, films, and virtual reality.	Luma , Poly
Product synthesis	Assistance with the design and development of physical products using spatial model methods in addition to generative design techniques.	Diabatix
Space synthesis	Automated design and modification of 3D spaces for use in virtual reality, gaming, and spatial analytics.	Common Sense Machines , Mirage , OpusAI



GENERATIVE AI TAXONOMY

Audio		
<p>Companies in this segment focus on the recording, manipulation, and playback of sound, speech, and music. Applications include human-machine interaction, music production, film and video production, podcasts, audiobooks, and more.</p> <p>Models: Whisper (OpenAI), data2vec (Meta), Jukebox (OpenAI)</p>		
Enhancement and editing	Techniques and tools used to improve the quality, clarity, and overall sound of recorded audio, as well as to edit, mix, and master audio tracks.	Alforithmic , Podcastle
Music generation	The use of generative algorithms to create new music, often based on user-specified parameters and musical styles.	LifeScore , Aimi , Splash
Summarization, search, and analysis	Tools for processing and analyzing large volumes of audio content to enable information retrieval and analytics.	Deepgram , Fathom
Synthetic voice	Computer-generated voices for various speech applications including dubbing, voice alteration, and voice replication.	Descript , Papercup
Text-to-speech/ speech-to-text	Automated conversion of written text to spoken audio, or spoken audio to written text, using natural language processing and machine learning algorithms. These techniques are used in a variety of applications, including accessibility tools, voice assistants, and automated transcription services.	AssemblyAI , WellSaid

Biotechnology
<p>Companies in this segment use artificial intelligence to analyze vast amounts of biological data, identify new targets, and optimize drug design. By harnessing the power of machine learning algorithms, they facilitate research and development of new pharmaceuticals, biologics, and other products. Their focus is on accelerating drug discovery and reducing costs associated with pre-clinical testing and clinical trials.</p> <p>Models: AlphaFold (Google/DeepMind), BioNeMo (Nvidia), ESM Metagenomic Atlas (Meta)</p> <p>Companies: Profluent Bio, Atomic AI</p>



GENERATIVE AI TAXONOMY

Code		
<i>Generative AI for creation, modification, or analysis of code. This includes developing advanced platforms that use AI algorithms to automate and streamline the coding process, improve code quality, and enhance the developer experience.</i>		
<i>Models: Codex (OpenAI)</i>		
Assistant	Tools that generate code and assist developers write more efficient or effective code.	GitHub , Replit
Documentation	Tools that generate documentation or assist developers in writing clear and concise documentation.	Mintlify
Web app building	Tools that generate code to build web applications or assist developers in building web applications.	Uizard

AI core		
<i>Companies that develop and provide tools and platforms to manage the model lifecycle, including data management, infrastructure, and end-to-end platforms for model development, deployment, and management.</i>		
Data	Platforms that provide tools for generating synthetic data, managing data pipelines, data storage, and data versioning.	Scale AI , Synthesis AI
Infrastructure	Tools and platforms that facilitate the creation and deployment of machine learning infrastructure, including containerization, orchestration, and scaling.	Banana , Anyscale
Platform	End-to-end platforms that enable the development, deployment, and management of machine learning models. These platforms often include features for data preparation, model training, and model deployment, as well as monitoring and management tools to ensure that models continue to perform effectively over time. This subsegment also includes research labs that focus on developing cutting-edge AI models.	OpenAI , Midjourney , Hugging Face



GENERATIVE AI TAXONOMY

Natural language interfaces		
<p>Companies in this segment specialize in creating software applications that enable users to interact with computers and devices using natural language. These applications use AI algorithms to process, understand, and generate human language in real time.</p> <p>Models: ChatGPT (OpenAI), GPT-4 (OpenAI), BLOOM, PaLM (Google)</p>		
Chatbots, voicebots, and personal assistants	This subsegment encompasses the creation of conversational agents for various purposes, including customer support, task automation, and information retrieval.	Asapp , Xembly , Vozy
General writing and editing	This subsegment includes tools that assist with grammar, syntax, and style, as well as those that help generate coherent text based on user input.	Grammarly , Flowrite
Marketing, sales, and copywriting	Companies in this subsegment offer tools for generating marketing and sales copy, including email campaigns, landing pages, sales communications, and social media posts.	Regie.ai , Lavender
Summarization and semantic search	This subsegment focuses on software that can extract key information from large volumes of text and provide summaries, as well as search for relevant information based on semantic understanding.	Neeva , Needl
Translation	Companies in this subsegment offer software tools for translating text from one language to another, including machine translation and human-assisted translation tools.	DeepL

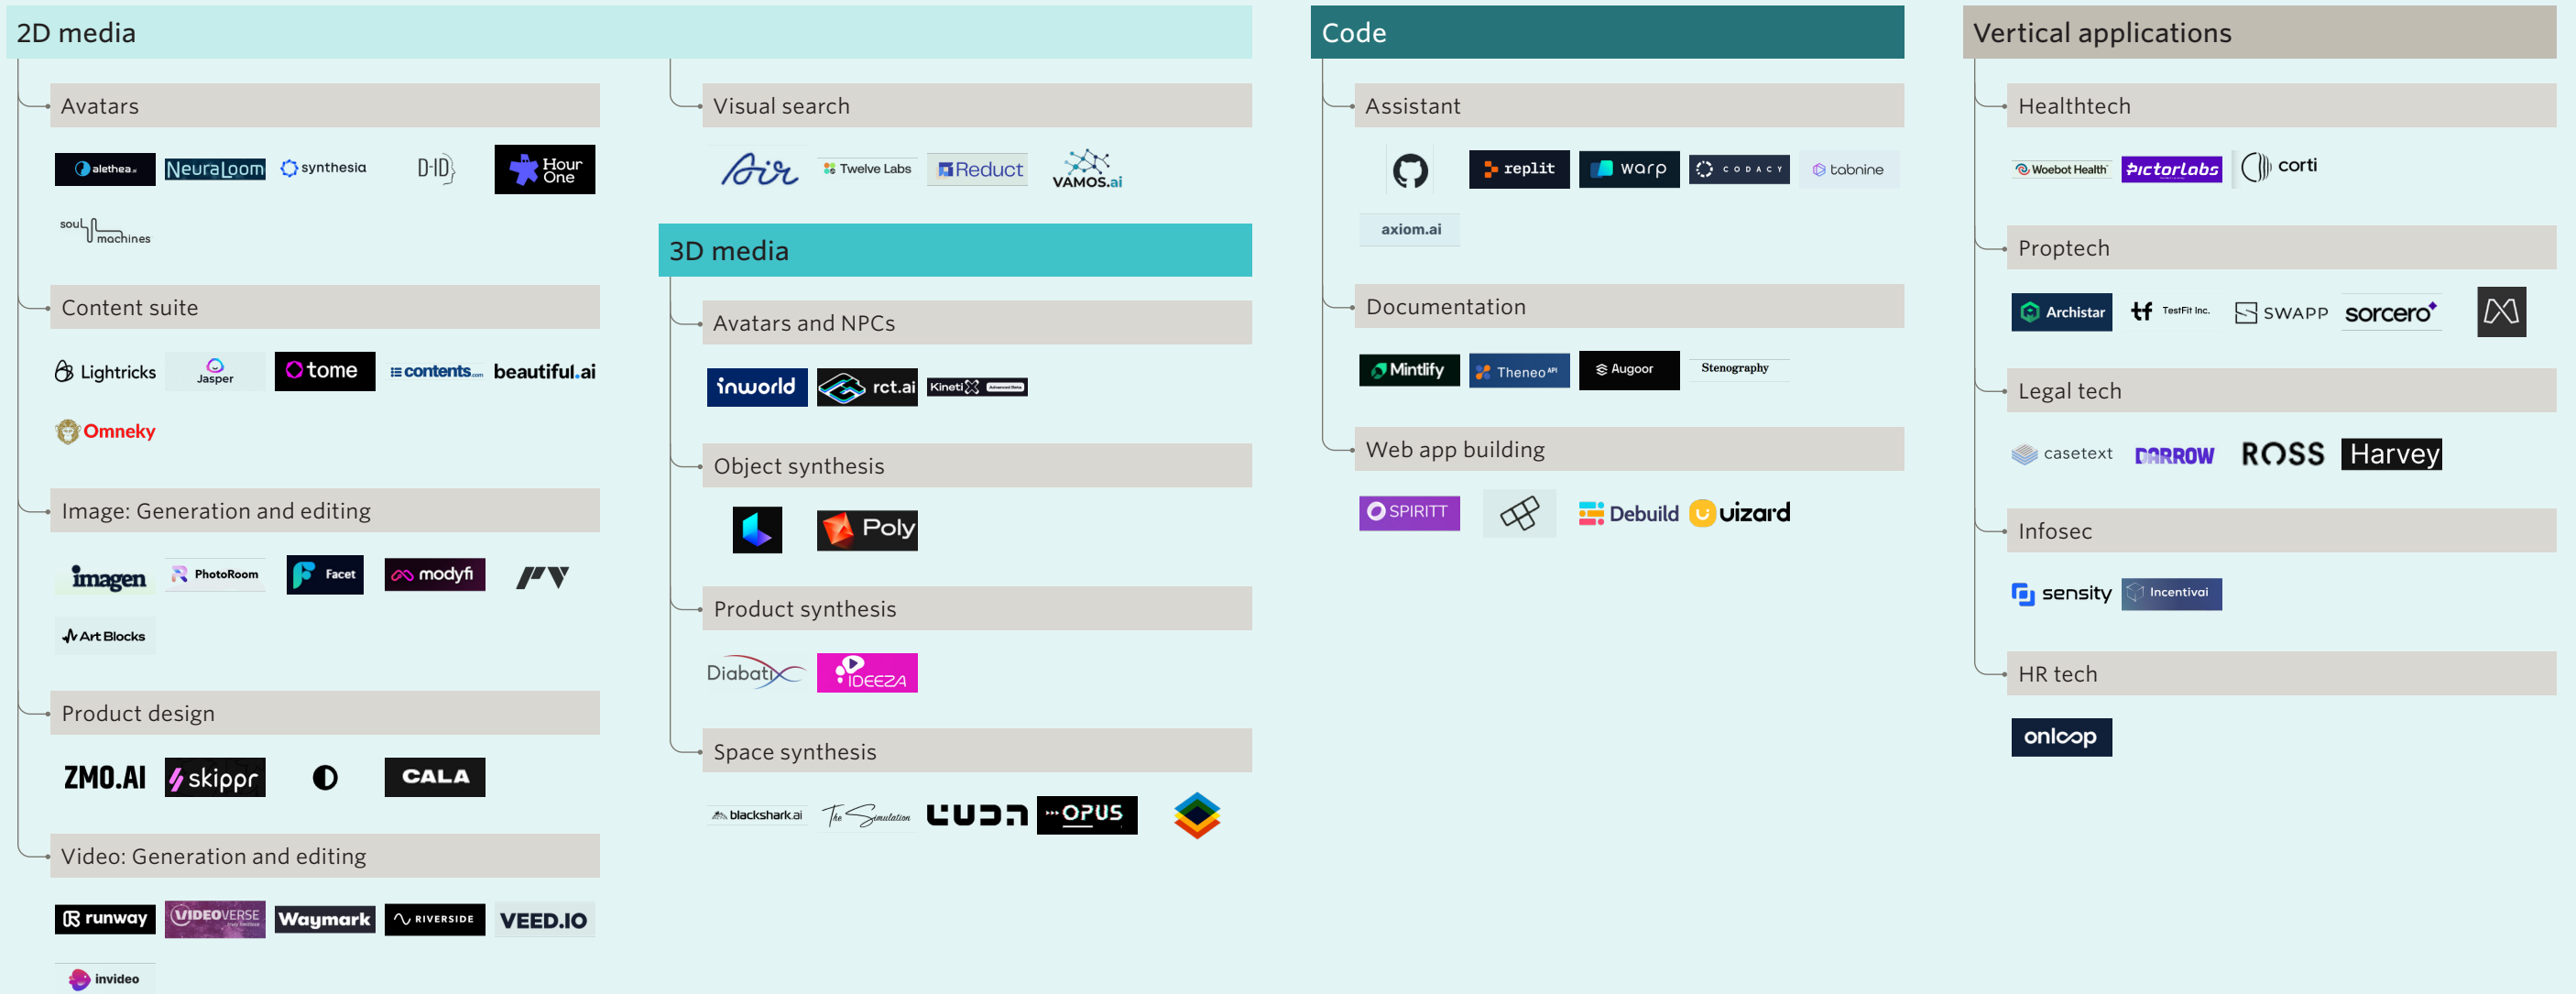
Vertical applications		
<p>Companies in this space develop generative AI solutions tailored to specific industries or use cases, rather than broadly applicable horizontal applications. Their software automates processes and improves outcomes for businesses within those industries or functions.</p>		
Healthtech	Companies that use AI to improve patient care and streamline administrative tasks.	Wobot Health , Pictor Labs
HR tech	Companies in this subsegment provide technology solutions for workforce management and human resources operations, including performance assessments.	ManageBetter
Infosec	Companies in this subsegment specialize in providing cybersecurity solutions, with a particular emphasis on detecting deep fakes.	Sensity
Legaltech	Companies that use AI to streamline legal processes, automate document analysis, and provide insights into case law and legal strategy.	Harvey , PatentPal
Proptech	Companies in this subsegment focus on providing technology solutions for real estate professionals, including property design, architecture, and real estate marketing.	Archistar , Maket , overwrite.ai



Generative AI 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.

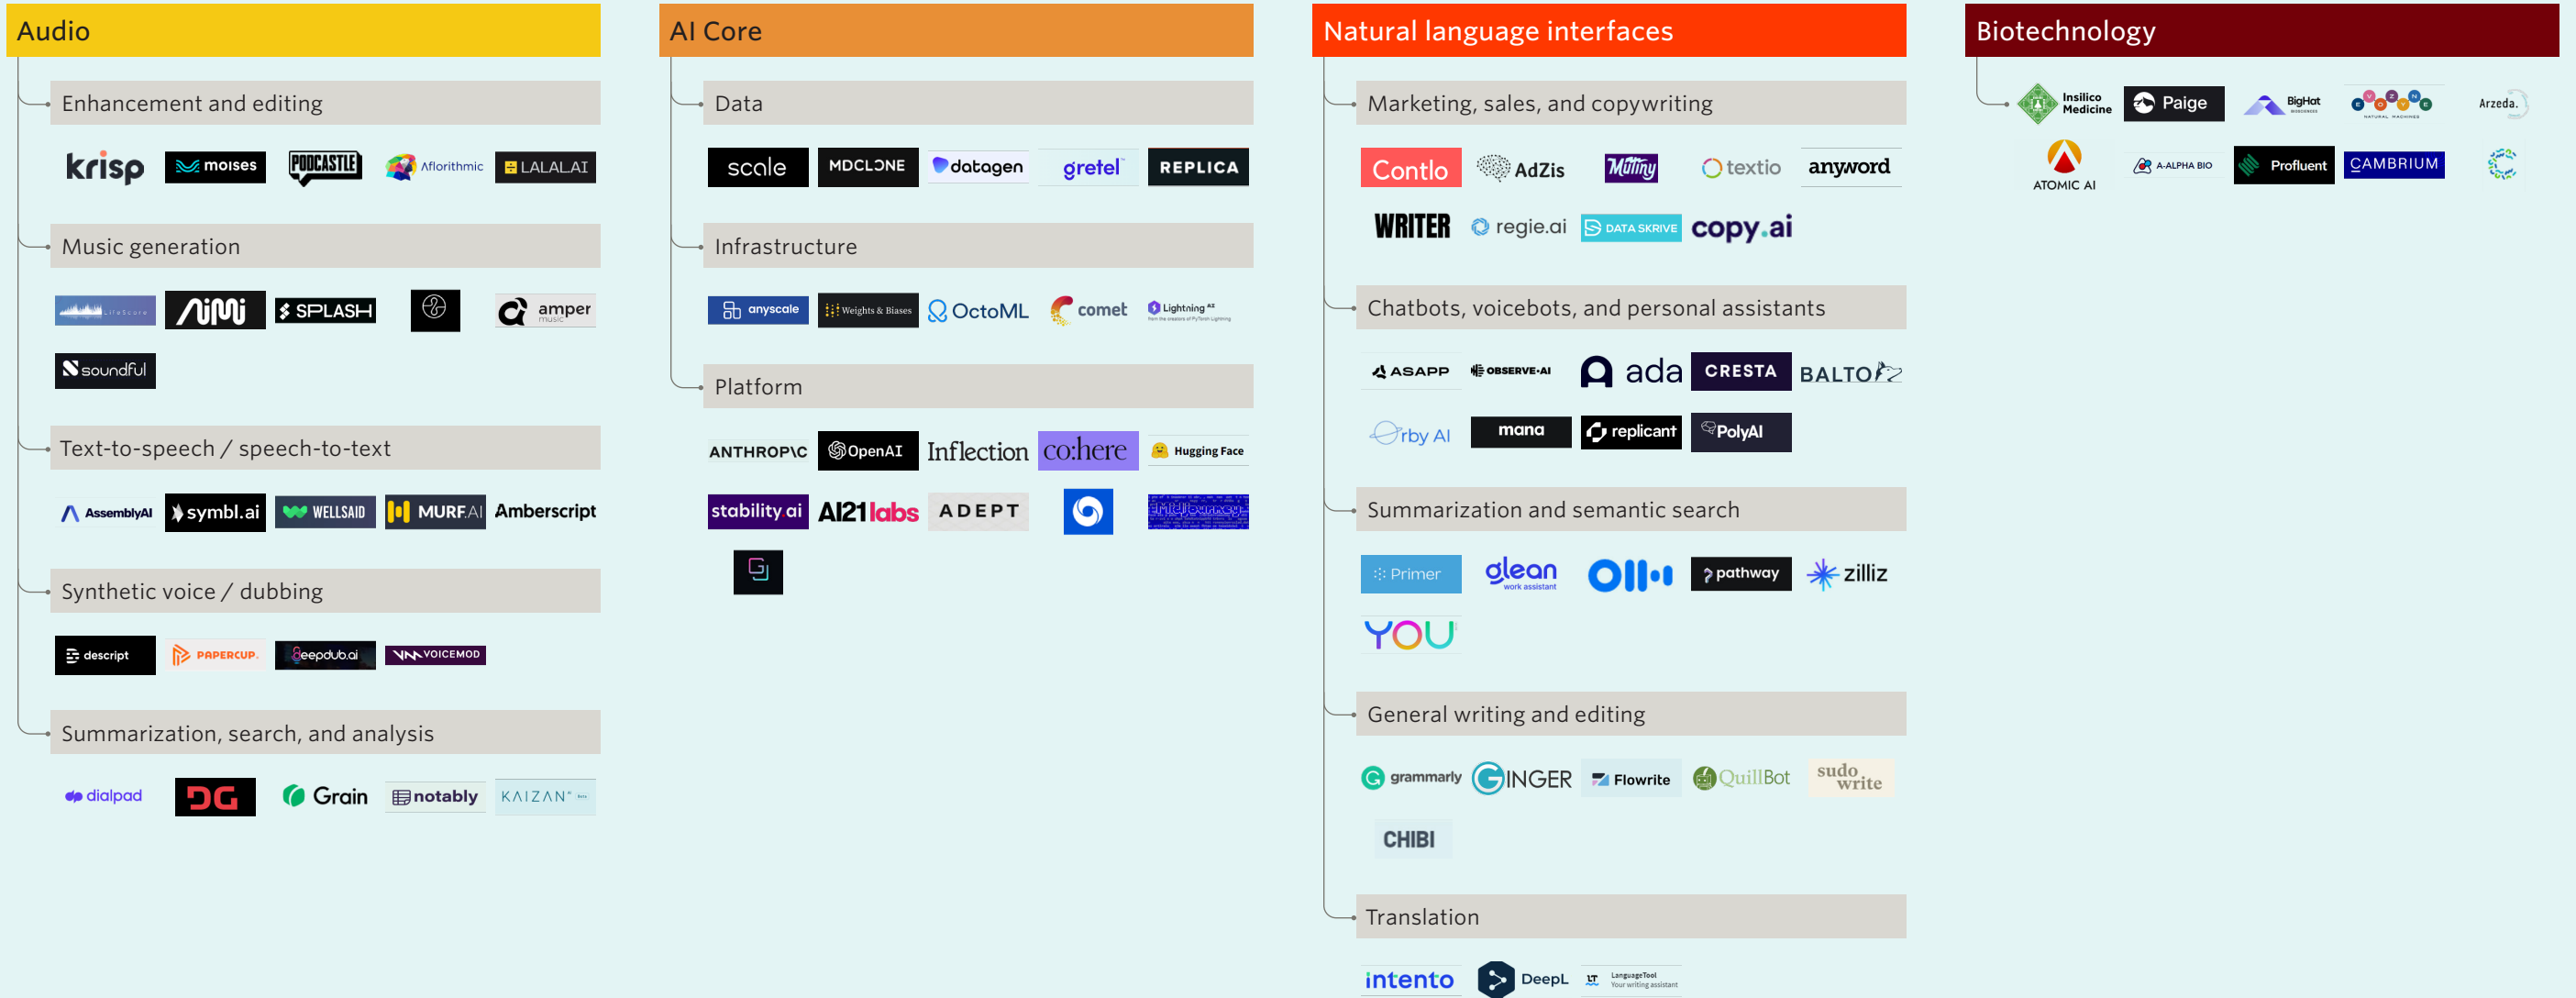




Generative AI market map


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.

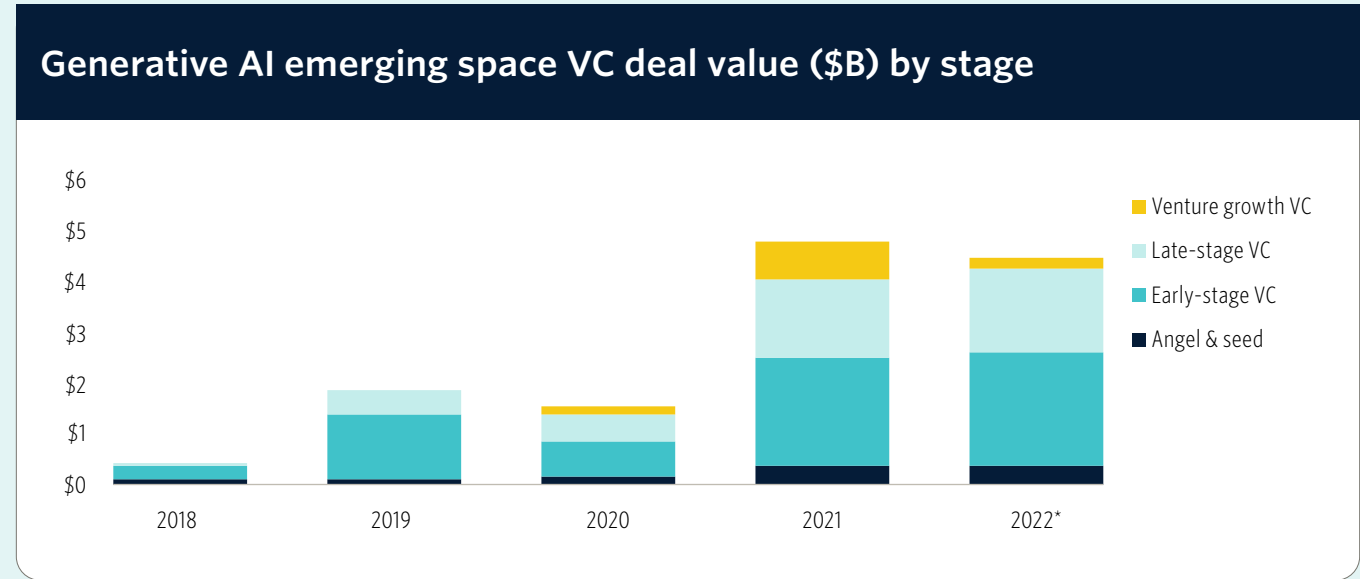





VC activity



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



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




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




VC ACTIVITY

VC deal count as share of total by segment 2018-2022*



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


VC deal value as share of total by segment 2018-2022*




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

2D media VC deal activity by use case



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

Audio VC deal activity by use case



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



Key players

Angel and seed key players*

Company	VC raised to date (\$M)	Segment	Category	City	Country
LifeScore	\$29.3	Audio	Music generation	Salisbury	UK
Ordaos	\$21.7	Biotechnology	N/A	New York	US
iGenius	\$20.0	Natural language interfaces	Chatbots, voicebots, and personal assistants	Milan	Italy
Twelve Labs	\$17.0	2D media	Visual search	San Francisco	US
Redesign Science	\$17.0	Biotechnology	N/A	New York	US
Tenyx	\$15.0	Natural language interfaces	Chatbots, voicebots, and personal assistants	Palo Alto	US
NLX	\$14.6	Natural language interfaces	Chatbots, voicebots, and personal assistants	New York	US
Mantium	\$14.5	AI core	Platform	Cincinnati	US
Peech	\$14.3	2D media	Video: generation and editing	Tel Aviv	Israel
Omneky	\$12.5	2D media	Content suite	San Francisco	US

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



KEY PLAYERS

Early-stage VC key players*

Company	VC raised to date (\$M)	Segment	Category	City	Country
Anthropic	\$1,304.0	AI core	Platform	San Francisco	US
OpenAI	\$1,010.1	AI core	Platform	San Francisco	US
Generate Biomedicines	\$420.0	Biotechnology	N/A	Cambridge	US
Inflection	\$265.0	AI core	Platform	Palo Alto	US
Cohere	\$170.0	AI core	Platform	Toronto	Canada
Jasper	\$146.6	2D media	Content suite	Austin	US
Woebot Health	\$129.9	Vertical	HealthTech	San Francisco	US
Enveda Biosciences	\$124.0	Biotechnology	N/A	Boulder	US
BigHat Biosciences	\$104.3	Biotechnology	N/A	San Carlos	US
Stability.AI	\$101.0	AI core	Platform	London	UK

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



KEY PLAYERS

Late-stage VC key players*

Company	VC raised to date (\$M)	Segment	Category	City	Country
Grammarly	\$400.0	Natural language interfaces	General writing and editing	San Francisco	US
GitHub	\$350.5	Code	Assistant	San Francisco	US
Lightricks	\$340.7	2D media	Content suite	Jerusalem	Israel
Asapp	\$312.6	Natural language interfaces	Chatbots, voicebots, and personal assistants	New York	US
Scale	\$277.7	AI core	Data	San Francisco	US
Anyscale	\$259.9	AI core	Infrastructure	San Francisco	US
Paige	\$220.0	Biotechnology	N/A	New York	US
Observe.AI	\$213.1	Natural language interfaces	Chatbots, voicebots, and personal assistants	San Francisco	US
Weights & Biases	\$200.0	AI core	Infrastructure	San Francisco	US
Ada	\$190.2	Natural language interfaces	Chatbots, voicebots, and personal assistants	Toronto	Canada

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



KEY PLAYERS

Top VC investors since 2018 by deal count*

Investor	Deal count	Angel and seed	Early-stage VC	Late-stage VC	Venture growth	Investor type
Andreessen Horowitz	21	5	9	5	2	VC
Tiger Global Management	20	1	9	8	2	VC
Sequoia Capital	20	5	12	3	0	VC
Amplify Partners	17	5	9	3	0	VC
Khosla Ventures	17	5	6	6	0	VC
Bloomberg Beta	15	4	8	3	0	Corporate VC
Index Ventures	15	4	9	1	1	VC
Soma Capital	15	11	3	1	0	VC
South Park Commons	14	5	6	3	0	VC
Alumni Ventures	14	7	5	2	0	VC

Source: PitchBook | Geography: Global | *As of March 10, 2023



Key partnerships

Alphabet/Anthropic

Alphabet's Google invested \$400.0 million in Anthropic and became the startup's preferred cloud provider through its Google Cloud Platform business. Google's internal language modeling innovation has lagged Meta and OpenAI, and the company has already leaned on partnerships with Cohere and other startups via corporate venture capital (CVC) deals after the company's acquisition of Deepmind in 2014. Anthropic is developing a competitive foundation model to OpenAI's GPT-3 and already surpasses all non-OpenAI foundation models in question answering accuracy, according to Stanford researchers' benchmarks.²⁶ We believe the company is valued around \$5 billion, granting a similar valuation per employee to OpenAI and Cohere.

AWS/Stability AI

Stability AI relies on a partnership with Amazon Web Services (AWS) for processing power. Stability AI CEO Emad Mostaque has acknowledged that while the company develops open-source AI, it has an exclusive cloud partnership with AWS.²⁷ To train the first version of the Stable Diffusion text-to-image model, the company self-funded the usage of 4,000 NVIDIA A100 graphics processing units (GPUs) from AWS to train computationally intensive diffusion models, according to its own disclosures.²⁸ We believe this partnership may increase in strategic relevance for AWS, given the limited progress the company has made in foundation model development to date.

²⁶: ["Holistic Evaluation of Language Models: Question Answering," Stanford Center for Research on Foundation Models, Percy Liang, Rishi Bommasani, et al. November 16, 2022.](#)

²⁷: ["Stability AI Selects AWS as Its Preferred Cloud Provider to Build Artificial Intelligence for the Future," Amazon, November 30, 2022.](#)

²⁸: ["How Stable Diffusion Was Built: Tips and Tricks to Train Large Models," AWS and Stability AI, Farshad Farahbakhshian, Emad Mostaque, et al., December 2, 2022.](#)

Microsoft/OpenAI

Microsoft's recent OpenAI-related product announcements and continued investment support for OpenAI show that the two will be tied at the hip going forward. OpenAI can serve as an R&D lab for Microsoft's future enterprise software products along with the giant's ambitions in search. OpenAI has incurred staggering computing costs for its ChatGPT preview, requiring additional commitments from its cloud partner to support commercial products such as ChatGPT Pro.

Recent rumors about deal terms for Microsoft's recent \$10.0 billion commitment suggest that Microsoft has bargaining power in this partnership to extract a guaranteed financial return and preferential licensing rights for OpenAI's foundation models.²⁹ These terms create disruption potential for startup AI research labs to define more favorable business models with open-source innovation. Startups with opportunities to benefit from this announcement include Anthropic AI, AI21 Labs, Adept, Cohere, and Stability AI.

²⁹: ["Microsoft Eyes \\$10 billion Bet on ChatGPT," Semafor, Liz Hoffman and Reed Albergotti, January 9, 2023.](#)



Key acquisitions/deals

Top VC deals by deal size*

Company	Close date	Deal size (\$M)	Pre-money valuation (\$M)	Post-money valuation (\$M)	Segment	Category	Lead investor(s)
OpenAI	July 2, 2019	\$1,000.0	N/A	N/A	AI core	Platform	Microsoft
Anthropic	April 29, 2022	\$580.0	\$3,420.0	\$4,000.0	AI core	Platform	Sam Bankman-Fried
Generate Biomedicines	November 18, 2021	\$370.0	\$1,080.0	\$1,450.0	Biotechnology	N/A	Altitude Life Science Ventures, ARCH Venture Partners, Fidelity Management & Research
Scale	April 13, 2021	\$325.0	\$6,975.0	\$7,300.0	AI core	Data	Dragoneer Investment Group, Greenoaks Capital Partners, Tiger Global Management
Inflection	April 28, 2022	\$265.0	\$960.0	\$1,225.0	AI core	Platform	Greylock Partners, Polaris Partners, The General Partnership
Insilico Medicine	June 1, 2021	\$255.0	\$375.0	\$630.0	Biotechnology	N/A	Warburg Pincus
Grammarly	November 18, 2021	\$200.0	\$12,800.0	\$13,000.0	Natural language interfaces	General writing and editing	Baillie Gifford
Anyscale	September 6, 2022	\$199.2	\$910.0	\$1,014.6	AI core	Infrastructure	Addition, Andreessen Horowitz, Intel Capital
Asapp	May 1, 2020	\$185.0	\$650.0	\$835.0	Natural language interfaces	Chatbots, voicebots, and personal assistants	Dragoneer Investment Group, Fidelity Management & Research
Dialpad	December 16, 2021	\$170.0	\$2,030.0	\$2,200.0	Audio	Summarization, search, and analysis	ICONIQ Growth

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



KEY ACQUISITIONS / DEALS

Top recent VC deals by deal size*

Company	Close date (2022)	Deal size (\$M)	Pre-money valuation (\$M)	Post-money valuation (\$M)	Segment	Category	Lead investor(s)
Anthropic	April 29	\$580.0	\$3,420.0	\$4,000.0	AI core	Platform	Sam Bankman-Fried
Inflection	April 28	\$265.0	\$960.0	\$1,225.0	AI core	Platform	Dragoneer Investment Group and Fidelity Management & Research
Anyscale	September 6	\$199.2	\$910.0	\$1,014.6	AI core	Infrastructure	Addition, Andreessen Horowitz, Intel Capital
Jasper	October 21	\$140.6	\$1,375.0	\$1,515.6	2D media	Content suite	Insight Partners
Observe.AI	April 12	\$125.0	\$700.0	\$825.0	Natural language interfaces	Chatbots, voicebots, and personal assistants	SoftBank Investment Advisers
Stability.AI	October 5	\$101.0	\$899.0	\$1,000.0	AI core	Platform	Coatue Management, Lightspeed Venture Partners, O'Shaughnessy Ventures
Glean	May 18	\$100.0	\$900.0	\$1,000.0	Natural language interfaces	Summarization and semantic search	Sequoia Capital
Hugging Face	April 6	\$100.0	\$1,900.0	\$2,000.0	AI core	Platform	Lux Capital
InstaDeep	January 25	\$100.0	N/A	N/A	AI core	Platform	Alpha Intelligence Capital, CDIB Capital Group
Insilico Medicine	June 1	\$95.0	\$800.0	\$895.0	Biotechnology	N/A	Prosperity7 Ventures

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



KEY ACQUISITIONS / DEALS

Jasper acquired Outwrite for an undisclosed amount on January 24, 2023.

Since its initial foray into copywriting generation, Jasper has become a generative AI leader. The company is expanding into a content suite. Outwrite adds AI writing suggestions including grammar, tone, and brevity. The acquisition primarily adds distribution given Outwrite's existing footprint of more than 1 million users. The move can signal a strategic shift toward direct-to-consumer generative AI services, enabling Jasper to benefit from the recent rush in demand for ChatGPT. Further, Outwrite's product is browser-native, enabling Jasper to more seamlessly integrate into user workflows and keep up with the experience of Microsoft Bing's integration in the Edge browser. The rapid evolution of this field requires acquisitions to keep pace.

Insider acquired MindBehind for an undisclosed amount on January 10, 2023.

E-commerce software leader Insider found value in MindBehind's AI chatbots. MindBehind turns social media platforms into virtual assistants through integrating conversational intelligence with common chat features such as Facebook Messenger, Instagram messages, and WhatsApp chat. The company recently integrated ChatGPT to personalize conversations, thereby demonstrating the competitive differentiation that startups can achieve in conversational AI by using the latest techniques. E-commerce companies will look to build on the momentum of text-based marketing via authentic chatbots.

Google acquired Alter for an estimated \$100 million on August 10, 2022.

Alter used computer vision to generate digital avatars similar to Apple's Memojis. The company previously raised only \$4.7 million, demonstrating the high strategic value of the company's technology and its viral adoption. The company was co-founded by former Big Tech product leaders who met via the Betaworks accelerator.³⁰ The app garnered significant traction in Apple's app store by optimizing avatars for diversity and inclusion. Precedent acquisitions included generative AI leader Roblox's acquisition of Loom.ai. This acquisition demonstrates that AI can accelerate consumer adoption of social media acquisitions and offers a precedent transaction for consumer-focused generative AI startups.

Spotify bought Sonantic for an estimated \$100 million on July 11, 2022.

Sonantic brought AI innovation to text-to-speech, earning two active US patents in speech processing and text-to-speech synthesis. Those patents enable the company to generate uniquely expressive voices that can be used in film and gaming. Spotify can use the technology for products such as its recently announced AI DJ, which lets one voice actor describe all the songs in its library as part of an immersive playlist. Generative AI applications in audio can achieve outstanding returns from traditional media companies.

³⁰: ["Facemoji, an App That Lets You Create Your Own Virtual Avatar, Has Over 2 Million Downloads by Tapping into What Makes Teens in 2019 Different From Other Generations," Business Insider, Paige Leskin, June 20, 2019.](#)



Glossary

Training: Training an AI model involves collecting and preparing a representative dataset, selecting an appropriate machine learning algorithm or neural network architecture, calculating the weights for each layer of a neural network to accurately model the dataset with matrix multiplications, evaluating its performance, and optimizing it to improve performance.

Machine learning approaches:

Supervised learning involves training a model using labeled data to accurately map input data to output labels. This approach is used to make predictions or classifications on new, unlabeled data.

Unsupervised learning is a machine learning approach that aims to discover hidden patterns or structures within unlabeled data. Unlike supervised learning, it does not rely on pre-existing knowledge of what the output should be.

Self-supervised learning is a type of unsupervised learning in which a model learns to predict missing features based on the existing features of a data point, with the objective being to learn a meaningful representation of the data.

Semi-supervised learning is a machine learning technique that combines labeled and unlabeled data to improve the performance of a model on labeled data. The model can learn more effectively by leveraging information from unlabeled data.

Reinforcement learning involves training a model to make decisions based on feedback from the environment. The goal is to learn a policy that maximizes a reward signal over time.

Transfer learning is a machine learning method that uses a model trained on one task as the foundation for a new, related task. The model can improve its performance on the target task by transferring knowledge learned on the source task.

Imitation learning involves a model learning to mimic a desired behavior from a teacher, such as a human expert. Eventually, the model will learn to reproduce the teacher's actions.

Reinforcement Learning from Human Feedback (RLHF) is a machine learning approach in which a model learns from feedback given by a human, rather than a reward signal from the environment. The goal is to learn a policy that maximizes the human's preferences or feedback.

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
Pharma Services