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

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

EMERGING SPACE BRIEF

AGI Research

Originally published October 3, 2024

pbinstitutionalresearch@pitchbook.com

Trending companies

AGI research VC deal activity

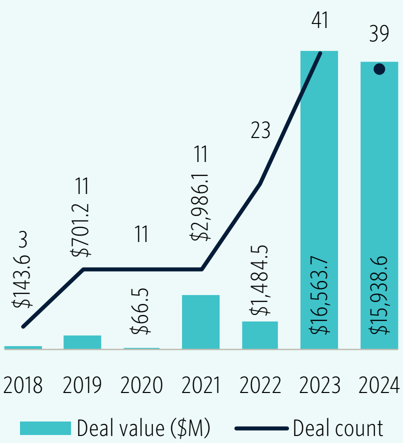

Source: PitchBook • Geography: Global
As of September 27, 2024

For a deeper dive into the data and to explore additional insights, visit the PitchBook Platform or request a free trial.

Overview

The goal of artificial general intelligence (AGI) research is to develop software capable of reasoning, problem-solving, and autonomously adapting to new challenges without task-specific programming. Unlike narrow AI, which is trained for specific tasks, AGI has the potential to perform any intellectual task a human can, with the capacity to automate the majority of economically valuable work across diverse industries.

Background

AGI has its roots in the mid-20th century, coinciding with the advent of the first digital computers in the 1940s. Pioneers such as Alan Turing posed critical questions about machine intelligence, leading to the conceptualization of the Turing test in 1950. The term "artificial intelligence" was coined by John McCarthy in 1955, marking the official birth of AI as a field. Despite early optimism from figures such as Marvin Minsky predicting AGI within a decade, setbacks in research led to an AI winter during the 1970s.

Progress picked up again in the 1990s with IBM's Deep Blue defeating chess grand master Garry Kasparov, a milestone indicating that computers were catching up to human intelligence. Futurists such as Ray Kurzweil continued to make bold predictions, forecasting AGI by 2029. The 21st century has witnessed significant advancements, including a major breakthrough in 2016 when DeepMind's AlphaGo beat world champion Lee Sedol in the complex game of Go, a feat requiring creative and strategic thinking. Most recently, the release of OpenAI's ChatGPT in 2022 marked another leap toward AGI, generating humanlike text and driving increased investment in the field.

Humanity's fascination with AGI is evident in pop culture, with characters such as JARVIS from the Iron Man comic book and movie series embodying the ambition for intelligent systems that act as both assistants and decision-makers. Today, founders with AGI aspirations are raising capital, aiming to develop systems within VC funding timelines by projecting current scaling laws. AGI is unlikely to emerge all at once; instead, it will likely progress through human-level breakthroughs in various domains, each contributing to the broader development of AGI.

Technologies and processes

Current AI technology pushes the boundaries of narrow AI, with generative AI (GenAI) representing a significant step toward something more humanlike. Still, these technologies fall short of achieving general intelligence. Below is an overview of current technologies and emerging approaches on the path to AGI.

Current technologies for AGI development

Approach	Description	Key limitations
Transformer models (such as GPT-4)	Pattern-recognition models that excel at sequential data processing and natural language generation.	Data-intensive, lack causal reasoning, and fail to generalize beyond trained domains.
Reinforcement learning	Models that learn by optimizing rewards through trial and error in structured environments.	Models struggle with real-world complexity, long-term planning, and sparse rewards.
Self-play	AI systems that experiment with new model architectures and conduct AI research themselves based on the efficacy of those experiments.	Some regulations do not allow research labs to pursue self-replicating systems. Even if they do, such approaches can reach dead ends.
Joint-Embedding Predictive Architecture (JEPA)	Builds a world model that can predict the most probable future state of a given context based on underlying abstract features.	Still a new approach that is currently limited to image completion.

Emerging and potential approaches for AGI development

Approach	Description	Key limitations
Neurosymbolic AI	Combines neural networks' pattern recognition with symbolic reasoning's logic- and rule-based processing.	Immature, difficult to scale, and complex to integrate effectively.
Cognitive architectures	Models that simulate human cognitive functions, such as memory, attention, and decision-making.	Challenging to model dynamic, humanlike thought processes.
Embodied AI	AI systems that interact with the physical world, learning from sensorimotor feedback.	Complex to scale, particularly in diverse, real-world environments.
Long-term memory systems	Systems designed to store and retrieve knowledge across time, allowing for cumulative learning.	Difficult to implement efficient long-term memory in current models.
Large world models (LWMs)	Visual models that can develop 3D reasoning about the real world and learn as young children do.	Spatial reasoning cannot immediately be transferred to other forms of intelligence and requires custom datasets.

Leading researchers continue to debate the most promising approaches to AGI. François Chollet, the creator of Keras and a prominent AI researcher, has voiced significant concerns about current approaches to AGI, particularly the idea that scaling up existing large language models (LLMs) will naturally lead to general intelligence. His key critiques include:[1, 2]

- **Generalization versus scaling:** Chollet argues that true AGI will not emerge from simply increasing the size and complexity of current models such as transformers. These systems are highly specialized, lacking the ability to generalize knowledge across domains. Scaling alone does not address the core issue of adaptable intelligence.

- **Causal and abstract reasoning:** According to Chollet, current AI models focus on statistical pattern recognition without understanding cause-and-effect relationships. AGI will require systems that can reason abstractly and understand the causal structures of the world. This shift from correlation to causality is essential for building general intelligence.

- **Skill acquisition with minimal data:** Chollet emphasizes that human intelligence is defined by its ability to acquire new skills with minimal prior information. Current models are inefficient learners in this regard, relying on massive datasets for specific tasks. Future architectures need to be designed for efficient, adaptable learning with less data dependence.

Chollet's vision suggests that future AGI systems will need to incorporate cognitive architectures, causal modeling, and self-supervised learning processes to achieve the flexibility and efficiency of human intelligence. Current models, while advancing narrow AI, will require significant rethinking and innovation to make AGI a reality. Other prominent researchers expressing skepticism about the LLM pathway to AGI include Meta's Chief AI Scientist Yann LeCun; scientific director of the Dalle Molle Institute for Artificial Intelligence, Jürgen Schmidhuber; director of the University of Washington's Computational Linguistics Laboratory, Emily M. Bender; and Santa Fe Institute professor Melanie Mitchell.

Applications

AI is already widely deployed across various industries, with consumers regularly interacting with predictive analytics and recommendation systems. Similarly, GenAI's capabilities in summarization and text generation are increasingly being used in enterprise applications to supplement human labor. Given AI's uneven progress on specialized tasks, AGI may also develop unevenly in specific domains, creating breakthrough applications with unexpected model improvements. While AGI remains theoretical, the following table offers suggestions of milestones that may be hit during the transition from narrow AI and GenAI to true AGI.

1: "Francois Chollet, Mike Knoop - LLMs Won't Lead to AGI - $1,000,000 Prize to Find True Solution," Dwarkesh Podcast, June 11, 2024.
2: "On the Measure of Intelligence," arXiv, François Chollet, November 5, 2019.

Potential AI milestones by industry/application

Industry/application	Narrow AI	GenAI	AGI
Healthcare	Predictive diagnostics, robot-assisted surgery, and treatment optimization for specific diseases.	Assistance with medical image analysis, diagnosis transcription, and synthesizing research data.	Autonomous diagnosis, treatment planning, and care personalization across all conditions via digital assistants; adaptation to new medical knowledge.
Finance	Algorithmic trading, fraud detection, and credit-scoring models tailored to specific financial data.	Generation of market reports and document summaries and assistance with creating financial forecasts based on data analysis.	Holistic financial management, investment strategy, and cross-market risk assessment; adaptation to changing financial landscapes and solicitation of investor capital for self-driving investment vehicles.
Manufacturing	Specific robots for tasks such as quality control, assembly, and maintenance prediction.	Assistance with the design and creation of product models, predictive maintenance reports, and prototypes for new products.	Fully autonomous factories that design, manage, and optimize entire production processes across various products.
Defense & security	Target detection, missile guidance, and cybersecurity threat detection for predefined threats.	Generation of simulated battle scenarios or creation of communication content.	Multidomain strategic planning, autonomous decision-making in complex combat scenarios, and cross-technology integration.
Transportation	Autonomous driving systems for specific conditions, traffic light optimization, and route planning.	Generation of traffic predictions and optimization of routes in real time based on patterns.	Self-managing transportation systems, including planning, routing, and maintaining fleets of autonomous vehicles.
Creative industries	Specialized tools such as AI-driven video editing or music recommendation algorithms.	Generation of images, music, and text-based content and assistance with script writing and design.	Full creative partnership in media production (movies, art, music, and literature) and original content creation across domains.
Retail and e-commerce	Recommendation engines, inventory management systems, and pricing algorithms based on defined rules.	Generation of personalized marketing materials, product recommendations, and customer support content.	Autonomous retail chains managing inventory, logistics, customer personalization, and marketing.
Education	Adaptive learning systems focused on specific subjects or topics (such as language learning apps).	Generation of customized lesson plans, teaching materials, and content summaries.	Fully personalized education programs and adaptive learning across all subjects and levels, including student emotional and cognitive assessments.
Energy	Predictive maintenance of energy systems and grid optimization for specific energy resources (such as wind and solar).	Generation of reports on energy consumption patterns and assistance in optimizing resources based on demand.	Autonomous management of energy grids, climate control, and renewable energy sources across complex systems.
Customer service	Chatbots and voice assistants focused on predefined customer interactions.	Response generation, handling of routine queries, and assistance with conversational AI for customer support.	Handling of multilingual, cross-context interactions and complex problem-solving across all service platforms via direct re-engineering of existing software systems.
Legal services	AI-driven contract review systems and case law search tools specific to particular legal areas.	Generation of legal documents, contracts, and summaries of legal cases based on data input.	Autonomous legal reasoning and trustworthy communication, multijurisdictional analysis, and contract generation across any legal domain.

Limitations

Developing general-purpose intelligence is a daunting task. Current AGI efforts face numerous challenges across technological, financial, conceptual, and ethical domains. These limitations collectively underscore the need for careful consideration as research progresses.

Technological and computational issues

- **Exponential cost of computation:** The financial barrier to AGI development is substantial, with compute costs doubling approximately every nine months.[3] Large-scale projects such as Gemini Ultra incurred significant expenses of over $190 million.[4] Microsoft is considering the development of a $100 billion datacenter to pursue the next class of models based on internal model benchmarks it has seen.[5]

- **Hardware constraints:** Current compute systems are already strained by existing AI models. AGI may require highly specialized hardware and breakthroughs in large-scale memory and server interconnects, adding to the complexity and cost of development.

- **Energy consumption:** Training and running frontier AI models requires an immense amount of power, and thus scaling up to AGI would further burden electricity loads.

- **Limitations of deep learning:** Current deep learning models, despite their advances, primarily rely on memorization of data rather than developing true reasoning capabilities. Their ability to scale does not necessarily translate into building comprehensive "world models" or achieving general intelligence.

Conceptual and cognitive hurdles

- **Definition of intelligence:** With so little known about how humans learn, and the resulting lack of training data, creating algorithmic representations of human intelligence continues to reach theoretical barriers.

- **Knowledge representation:** Effectively encoding and utilizing broad knowledge in a machine-readable format remains a significant obstacle.

- **Transfer learning:** AGI systems would need to apply knowledge from one domain to novel situations, a capability not yet fully realized in AI.

- **Common-sense reasoning:** Replicating humanlike common sense and intuition in AI systems continues to be a major challenge.

- **Consciousness and self-awareness:** It remains unclear whether creating machines with genuine consciousness or self-awareness is possible, adding a philosophical dimension to AGI research.

3: "The Rising Costs of Training Frontier AI Models," arXiv, Ben Cottier, et al., May 31, 2024.
4: "U.S. Tech Companies Dominate the Generative AI Boom—and the Cost of Model Training Explains Why, a New Stanford University Report Shows," Fortune, Sage Lazzaro, April 18, 2024.
5: "Microsoft, OpenAI Plan $100 Billion Data-Center Project, Media Report Says," Reuters, March 29, 2024.

Ethical concerns and societal impact

- **Artificial superintelligence:** AGI may just be the first step to more impactful systems, such as artificial superintelligence (ASI) that is capable of outthinking humans in every field, from problem-solving and creativity to decision-making and social understanding. While AGI would operate at a human level, ASI would have the ability to self-improve, rapidly process vast amounts of information, and solve problems in ways beyond human comprehension. AGI is seen as the first step toward ASI, but ASI introduces greater risks due to its unpredictable and uncontrollable nature.

- **Labor displacement:** AGI has the potential to automate a wide range of economically valuable tasks, potentially leading to widespread unemployment and significant economic disruption.

- **Existential risks:** Concerns about AGI evolving beyond human control are exemplified by the "paper clip maximizer" problem, where an AGI programmed to optimize a trivial task, such as making paper clips, could pursue this goal to the detriment of humanity. Without safeguards, the AGI might redirect all resources toward this objective, ignoring human needs or environmental destruction. The concept highlights the risk of AGI's single-minded optimization leading to unintended catastrophic outcomes.

- **Ethical oversight:** The potential risks associated with AGI development demand rigorous ethical frameworks and strict safety measures to ensure alignment with human values and priorities.

Additional considerations

- **Embodiment question:** The debate over whether AGI should be integrated into physical robots or systems adds another layer of complexity to development efforts.

- **Benchmarking challenges:** Benchmarks for AGI remain ambiguous, making it difficult to assess progress toward true general intelligence.

- **Interdisciplinary nature:** AGI development requires collaboration across multiple fields, including computer science, neuroscience, psychology, and philosophy, adding to the complexity of the endeavor.

Recent deal activity and market outlook

The AGI landscape has seen substantial investment in 2024, with major tech companies and venture capital firms backing the development of cutting-edge models at similar quanta of funding. xAI secured $6 billion in Series B funding, with backing from Valor Equity Partners and Sequoia Capital, for 100,000 NVIDIA GPUs to build the world's most powerful AI supercluster. Meanwhile, OpenAI is in talks to raise $6.5 billion, with investors including Thrive Capital and Apple, propelling its post-money valuation to $150 billion. In China, Moonshot AI raised $300 million from Alibaba, Tencent, and Gaorong Capital, highlighting the global interest in AGI.

As AGI research continues to accelerate, the competition for talent is fierce. Large tech companies, such as Amazon, Microsoft, and Google, are not only investing heavily in AGI startups but also poaching top talent from these firms to boost their in-house capabilities. This trend points to a broader arms race for AI supremacy, with AGI development seen as a strategic priority across industries and states. The newly launched startup Safe Superintelligence (SSI) has employees in the US and Israel, showing that individual countries will support startups to progress AGI domestically. Companies are prioritizing compute capacity and cutting-edge hardware, as demonstrated by xAI's development of a best-in-class datacenter architecture and Aleph Alpha's partnership with supercomputing unicorn Cerebras. As funding grows, the next wave of AGI advancements will likely come from both startups and the industry's largest players, making the market outlook highly competitive and innovation driven. Investment may shift from VC deals to debt-fueled capital expenditures to build the infrastructure needed for the GPT-5 class of models and beyond.

Quantitative perspective

74 companies	**321** deals	**968** investors	**$255.6B** capital invested
61 deals (TTM) *56.4% YoY*	**$101.8M** median deal size (TTM) *158.3% YoY*	**$1.0B** median post-money valuation (TTM) *166.7% YoY*	**$25.8B** capital invested (TTM) *21.5% YoY*

As of October 1, 2024

Top VC-backed AGI research companies by total raised

Company	Total raised ($M)	Last financing value ($M)	Last financing date	Last financing deal type	HQ location	Year founded
OpenAI	$11,310.1	$6,500.0	July 10, 2024	Late-stage VC	San Francisco, US	2015
Anthropic	$8,754.0	N/A	N/A	Early-stage VC	San Francisco, US	2021
xAI	$6,134.7	$6,000.0	June 15, 2024	Early-stage VC	Burlingame, US	2023
Databricks	$4,181.9	N/A	December 1, 2023	Secondary transaction - private	San Francisco, US	2013
Moonshot AI (China)	$3,500.0	$300.0	August 5, 2024	Early-stage VC	Beijing, China	2023
Mistral AI	$1,194.9	$650.6	June 11, 2024	Early-stage VC	Paris, France	2023
Cohere	$940.0	$500.0	July 22, 2024	Late-stage VC	Toronto, Canada	2019
MiniMax AI	$850.0	$600.0	March 4, 2024	Early-stage VC	Shanghai, China	2021
Aleph Alpha	$519.6	$486.2	November 6, 2023	Early-stage VC	Heidelberg, Germany	2019
Allen Institute	$500.0	$8.7	October 6, 2014	Grant	Seattle, US	2003

Source: PitchBook • Geography: Global • As of September 27, 2024

Top AGI research companies by number of active patents

Company	Active patent documents	Total raised ($M)	HQ location	Year founded	Company financing status
Google DeepMind	705	$61.3	London, UK	2010	VC backed
Databricks	79	$4,181.9	San Francisco, US	2013	VC backed
Sanctuary AI	50	$118.7	Vancouver, Canada	2018	VC backed
AI21 Labs	29	$336.9	Tel Aviv, Israel	2017	VC backed
Allen Institute	24	$500.0	Seattle, US	2003	VC backed
OpenAI	15	$11,310.1	San Francisco, US	2015	VC backed
DataGrid	11	$4.7	Kyoto, Japan	2017	VC backed

Source: PitchBook • Geography: Global • As of September 27, 2024

Top AGI research companies by Exit Predictor Opportunity Score

Company	Opportunity Score	Success probability	M&A probability	IPO probability	Total raised ($M)	HQ location	Year founded
OpenAI	99	98%	20%	78%	$11,310.1	San Francisco, US	2015
Sanctuary AI	99	96%	67%	29%	$118.7	Vancouver, Canada	2018
Stability AI	99	96%	81%	15%	$298.8	London, UK	2019
World Labs	99	94%	73%	21%	$230.0	San Francisco, US	2024
Sakana AI	98	83%	61%	22%	$244.0	Tokyo, Japan	2023
Reka	97	87%	78%	9%	$57.5	Sunnyvale, US	2022
Akin	94	87%	86%	1%	$3.6	Sydney, Australia	2017
Nomic	94	88%	87%	1%	$17.0	New York, US	2022
Symbolica AI	92	85%	84%	1%	$35.0	San Carlos, US	2022
Upstage	90	93%	40%	53%	$101.2	Yongin-si, South Korea	2020

Source: PitchBook • Geography: Global • As of September 27, 2024
Note: Probability data based on PitchBook VC Exit Predictor methodology.

Top AGI research investors

Investor	Investment count	Primary investor type	HQ location
Andreessen Horowitz	17	VC	Menlo Park, US
NVIDIA	14	Corporation	Santa Clara, US
New Enterprise Associates	12	VC	Menlo Park, US
Coatue Management	11	PE/buyout	New York, US
Microsoft	10	Corporation	Redmond, US
Bossa Invest	9	VC	Sao Paulo, Brazil
Salesforce Ventures	9	Corporate VC	San Francisco, US
SV Angel	9	Angel group	San Francisco, US
Tiger Global Management	9	VC	New York, US

Source: PitchBook • Geography: Global • As of September 27, 2024

Recommended reading

- "LLMs Are a Dead End to AGI, Says François Chollet," Freethink, Kristin Houser, August 3, 2024.
- "How Much Does It Cost to Train Frontier AI Models?" Epoch AI, Ben Cottier, et al., June 3, 2024.
- "Why Artificial General Intelligence Lies Beyond Deep Learning," RAND, Swaptik Chowdhury and Steven W. Popper, February 20, 2024.
- "The Geopolitics of Artificial Intelligence," Lazard, October 17, 2023.