Policy Brief

China's Cognitive AI Research

Emulating Human Cognition on the Way to General Purpose AI

Authors
William C. Hannas
Huey-Meei Chang
Max Riesenhuber
Daniel H. Chou

CSET CENTER for SECURITY and EMERGING TECHNOLOGY

July 2023

Executive Summary

China's intent to create broadly capable artificial intelligence, also called "artificial general intelligence" (AGI), was announced in its 2017 "New Generation AI Development Plan" and is championed by leading Chinese scientists and AI institutions.

This study assesses the plausibility of these claims by examining Chinese scientific papers published in Chinese and English between 2018 and 2022 for evidence of related research. While most such papers are on routine AI applications, a significant body of research was found on AGI precursor technologies, indicating that China's claims to be working toward artificial general intelligence are genuine and must be taken seriously.

The study reaches the following conclusions:

- Published scientific studies indicate China is actively researching general AI.
- Chinese research on advanced (general) AI is shared over a broad talent base.
- The greatest concentration of Chinese AGI research is in the Beijing area.
- Global contributions support the research but are not its main drivers.

Given the need to monitor AI developments into the future, this study outlines the methods it used to reach these conclusions. These local efforts, a demonstration of open sources' ability to track foreign science and technology (S&T) in general, cannot be sustained independently and should be augmented by a national program to avoid surprise and ensure access to global scientific publications.

Table of Contents

Executive Summary 1
Introduction 3
Methodology 4
Mapping the Target 4
Document Retrieval 5
Analytic Findings 7
Caveats and Dependencies 15
Recommendations 16
Authors 18
Acknowledgments 18
Endnotes 19

Introduction

As artificial intelligence (AI) moves from narrow to broad applications, developers worldwide see benefit in emulating advanced—but computationally elusive—features of human intelligence. These features include transfer learning, reasoning, creativity, intuitive understanding (common sense), and others described below, along with peripheral aspects of intelligence such as affect and motivation. Since humans perform these functions natively, the assumption, shared by Chinese researchers, is that modeling aspects of human cognition will lead to more capable AI.[1]

This paper examines Chinese approaches to these challenges. Prior studies identified China-affiliated scientists pursuing "General Purpose AI" (通用人工智能) based on their public claims, many of whom seek clues from brain science to overcome current bottlenecks in machine learning.[2] The present study delves deeper by examining their *actual research*.

China's 2017 "New Generation AI Development Plan" names General AI as a key to dominating world AI and a first-mover advantage, while asserting the value of brain-inspired (类脑) AI as a path toward General AI.[3] A link between Chinese artificial and human intelligence research is further evidenced in China's declared aim to "merge" (混合) AI and neuroscience, and in National Natural Science Foundation of China funding categories that support this research.[4] The present study examines how these aspirations are reflected in the published scientific record.

Beyond its topical goals, this study also introduces a procedure for analyzing foreign scientific publications. The value of open-source information—long known to China, which has been using OSINT for decades[5]—is increasingly appreciated within the U.S. and allied governments as a vital supplement to secrets-based assessments of foreign science.[6] Practices used in this present study of China's AI are extensible to other topics and countries and worth describing in detail.

China's cognitive AI research will improve its ability to field robots, make smarter and quicker decisions, accelerate innovation, run influence operations, and perform other high-level functions reliably with greater autonomy and less computational cost, elevating global AI risk and the strategic challenge to other nations. AI is also an *enabling technology*, whose greater value is realized when applied to other disciplines.[7]

Methodology

This paper follows an open-source exploitation model for S&T studies developed by the authors and their circle of colleagues. There are eight steps:

1. Generate Chinese and English search terms that map the discipline.
2. Run these terms on the internet and through academic journal holdings.
3. Downselect the returned papers based on expert assessment of their content.
4. Use the metadata of selected papers to build new (internet) search queries.[8]
5. Amalgamate the corpora and perform entity resolution on authors and affiliations.
6. Assess the nature and quality of the content through expert analysis.
7. Determine key players and institutions based on content and authorship patterns.
8. Make analytic judgments about the field and where things are headed.

These stages and their findings are described in the following sections.

Mapping the Target

The authors identified several so-called AI "bottlenecks" from prior research, the AGI literature, and expert consultations, seen as potentially amenable to brain-inspired solutions or as cognitive capabilities that improved machine learning (ML) algorithms might address. These targets were sorted into a dozen categories as follows:

- affective computing (artificial emotion, sentiment analysis from facial images, body language, and text)
- applications (with human in the loop: brain-computer interfaces (BCI), neuromodulation, augmented cognition, human-computer interaction)[9]
- awareness (self-consciousness, first-person subjectivity, theory of mind, empathy, moral reasoning, cognitive control)
- BI-AI (brain-inspired AI, if more than perception, including cognition and "theories of everything")
- embodiment (navigation, robotics as a precondition for environmental interaction and learning)
- generative (including large language models and image generation models, although not a specific focus of this study)
- hardware (memristors, neuromorphic chips, processing-in-memory, artificial neurons, spike-based computing)

- memory (associative, episodic—cognitive recall internally generated or as responses to environmental stimuli—not computer "memory" requirements)
- pattern recognition learning (video, imitation, transfer, continuous, autonomous, zero/one/few-shot, and meta-learning, small model)
- perception (vision, adaptive attention, situational awareness, navigation with a primary focus on application in vision)
- reinforcement learning/imagination (planning, creativity, strategies/games, reward-based learning, navigation, control)
- sensemaking (intuition, reasoning, causation, uncertainty, inference, abduction, commonsense, concepts, language, neural basis of intelligence)

Papers by Chinese scientists that discuss AGI as an abstract or philosophical goal but do not treat specific bottleneck areas were also collected.

Bilingual search terms were created for each of these topics and their descriptors. In some cases, they were the topic terms themselves; in other cases related concepts that point to the topic. If the base term was in English, Chinese translations were tested for authenticity, and vice-versa.

The terms were run against the CSET merged corpus of scholarly literature including Digital Science Dimensions, Clarivate's Web of Science, Microsoft Academic Graph, China National Knowledge Infrastructure, arXiv, and Papers With Code; CNKI via commercial subscription;[10] Chinese S&T journal metadata collections built independently; and the internet. Additional searches were done using keywords specified in the documents retrieved from the initial search.

Queries were also run using the terms "brain-inspired AI" / 类脑人工智能 and "general purpose AI" / 通用人工智能 with the names of China's major journal aggregators.[11]

Document Retrieval

Some 500 documents were retrieved for years 2018-2022.[12] As expected, many describe ML research aimed at cognitive capabilities—studies that are brain-inspired in a de facto sense. The decision was made to treat these documents as within scope.[13]

Conversely, documents that reflect pure neuroscience without a significant AI component, or studies analyzing neural data with the help of AI, were stricken from the pool. Many documents in the "applications (human-in-the-loop)" category on brain-computer interfaces (BCI) were consigned to a CSET follow up study.[14]

The surviving records were divided into “interesting” (about one-third of the corpus) and “less interesting” based on (1) a subject matter expert’s (SME) review, (2) the relative standing of the journal or forum, (3) status of the funding body, (4) citation frequency, and (5) the number of CNKI downloads.[15] The aim was to identify papers that best exemplify the study’s goals and use them to expand the corpus by searching their authors' names for papers missed in the initial search. Some 130 papers were retrieved this way, bringing the count to 506, after the reductions described above.

Additional searches were then done based on:

- top Chinese AI scientists identified in prior CSET research;
- the CNKI classification codes assigned to selected documents;
- CNKI’s compilation of “related papers” for selected documents;
- the complete record of 人工智能 (*Artificial Intelligence*) as a check on the above;[16]
- National Natural Science Foundation of China funding codes;
- featured studies on the websites of Chinese AI research institutes;
- English language academic publications by Chinese scientists.

This last step, counterintuitively, was the most challenging because of the need to reconstruct the authors’ Chinese names and affiliations.[17] The bibliographies of Chinese language papers were the source of many such English language studies. Others were chosen from:

- NeurIPS and AAAI conference papers for the years 2021-2022;[18]
- *Frontiers of Information Technology and Electronic Engineering*;[19]
- *International Journal of Automation and Computing*, and its successor:
- *Machine Intelligence Research*.[20]

This hands-on strategy reflects the difficulty of programming searches for a topic that is intrinsically ill-suited to measurement. Namely, AI is not a well-defined discipline, “advanced” AI less so, and there is no clear line between innovative and run-of-the-mill progress, absolutely or relative to China’s vision of where the enterprise is headed.

The outcome of these searches and several rounds of pruning for relevancy was a spreadsheet of curated metadata and the full texts of 850 studies by China-affiliated AI scientists on topics judged to be precursors to AGI. This data, observations made during their compilation, and material from additional research form the basis for the following analytic judgments.

Analytic Findings

1. ***Published scientific studies indicate China is moving toward general AI.***

A prior CSET study cited claims by Chinese scientists that China is moving toward general artificial intelligence.[21] That study took into account both machine learning, including large language models, and other approaches based on cognitive and neuro-sciences. A list of AGI "indicators" and keywords was appended to the study to support later investigations.

Based on those indicators and the counsel of subject matter experts, a corpus of 850 papers authored by China-affiliated scientists was retrieved, whose content met thresholds for advanced AI, understood as "brain-inspired" literally, and derivatively in the sense of computational efforts to emulate human cognitive skills without reference to biological processes. Since these skills are interdependent, 105—12 percent—of the 850 Chinese papers were assigned to two or more of the topical categories described above for a total of 995 data points.

Figure 1. Distribution by Topic of Chinese Academic Papers on Advanced AI

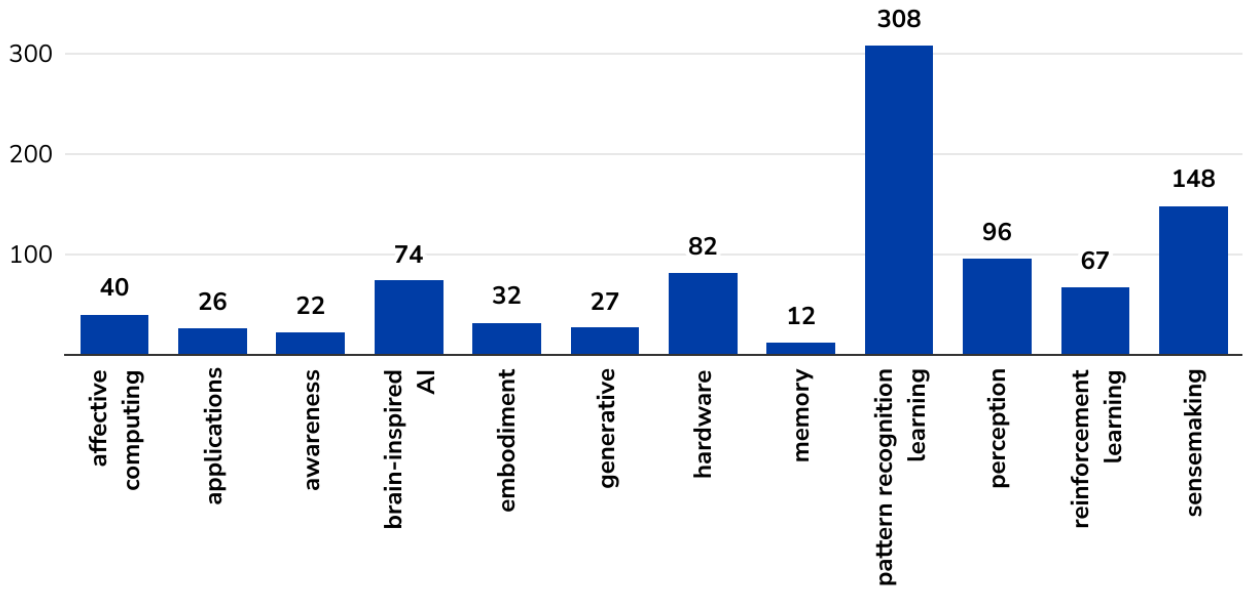

Source: Multiple (as described above)

Figure 1 plots the distribution of papers by content. An additional category, studies on AGI that do not address specific problems, attributes, or topical areas, has 61 papers, suggesting the end state of these precursor studies is also an area of Chinese scholarly

interest. While the variance in size between topics is worth exploring, the main point is that the data, as a whole, demonstrates that Chinese scientists are researching all areas identified in principle as supporting general AI.

This assessment is consistent with the results of an earlier study (2020) based on a smaller machine-generated sample of 352 Chinese academic papers covering 10 topical areas indicative of a brain-inspired (BI) AI program.[22] Documents in the 2020 study have a mean publication date of 2016. Those in the present study date on average from 2021. Although the categories used in the two studies differ in detail, both focus on attributes associated with AGI projects.

Since the present corpus was hand curated, pre-filtering was possible under SME guidance. Within the total pool of discoverable papers, the number of Chinese AI papers on basic science and advanced applications—this study's target—was overshadowed by routine applications-oriented research, a known phenomenon.[23] Most of those papers were excluded from the study.

These off-topic papers, eliminated by design, numbered in the thousands and centered on self-driving automobiles, pilotless drones, safe cities, pedagogy, etc., whose technology did not rise to the "advanced" level elicited. In addition, there was an unusually large number of papers on facial, gait, and emotion recognition; sentiment analysis of visual and auditory media; (errant) behavior prediction; and military applications that were bypassed because the technology was old and the motivation more political than exploratory. Where the technology was advanced or aimed at more general use, samples of this genre were also included—and are responsible in part for the spike in "pattern recognition learning."

Hence the corpus does not—and is not meant to—represent Chinese AI research in general but rather as substantiation of Chinese claims, made in other open-source venues, that the country is building toward the "first-mover advantage" and "hybrid human-machine" goals articulated in its 2017 "New Generation AI Development Plan."[24] China's representation in top AI publishing venues, the volume of on-topic research in credible domestic journals, and our own assessment of the papers' exploratory depth show that the effort is genuine and should be taken seriously.

2. ***Chinese research on advanced (general) AI is shared over a broad talent base.***

The corpus of 850 documents has more than 2,000 unique authors and co-authors, which suggests that Chinese interest in aspects of general AI is widespread. Each document has one to a dozen or more (co) authors, but typically just three to five. It was infeasible to capture this level of detail entirely in a working spreadsheet—and unnecessary given this study's limited goal of validating claims about the nature of China's advanced AI research.

Accordingly, the following rules were used for author extraction: three, or exceptionally four, authors were selected per paper, in this order: first author, second author if the work was shared equally, corresponding author(s) if specified, and last author (often the senior person overseeing the study).

Some 1,760 unique Chinese authors were extracted (foreign input is discussed below) to the study's spreadsheet. Among them, 245 are on the bylines of two or more papers, 32 have five or more papers, and the following authors, highlighted below, each have seven or more papers:[25]

- GAO Wen (高文). Gao is a professor in computer science (CS) at Peking University (PKU) and director of the Peng Cheng Laboratory. He has more than 1,000 papers in pattern recognition, computer vision, and object-oriented video coding. Gao has presided over some 20 national level projects.[26]
- HUANG Tiejun (黄铁军). Huang is vice dean of PKU's Institute for Artificial Intelligence, dean of the Beijing Academy of Artificial Intelligence, and is on China's AI governance board. His specialties are neuromorphic computing and vision processing.[27]
- LIU Jianwei (刘建伟). Liu is a professor at the China University of Petroleum and senior member of the Chinese Association for Artificial Intelligence. Liu researches machine learning, pattern recognition, and non-linear systems. He has published 260 papers.[28]
- WANG Feiyue (王飞跃). Wang holds high posts at the Chinese Academy of Sciences Institute of Automation (CASIA) and the National University of Defense Technology (NUDT). He has extensive overseas experience, was an early "talents" selectee, and a recipient of many high-level awards. He researches social computing, knowledge automation, and intelligent control.[29]
- WU Fei (吴飞). Wu is a professor and vice dean of Zhejiang University's College of CS, director of the university's Institute of AI, and on multiple editorial and AI governance boards. His research interests include ML and multimedia analysis and retrieval.[30]

- XU Bo (徐波). Xu is CASIA's director, dean of the Chinese Academy of Sciences (CAS) University's School of Artificial Intelligence, and holds leadership positions in several other top AI research and governance organizations. He researches BI-AI and connectomics (brain mapping).[31]
- YANG Yaodong (杨耀东). Yang researches reinforcement learning and multi-agent systems at the Beijing Institute for General Artificial Intelligence (BIGAI) and PKU's Institute for Artificial Intelligence. He was a principal research scientist at Huawei, UK, and an assistant professor at King's College London.[32]
- ZENG Yi (曾毅). Zeng is deputy director of CASIA's Research Center for Brain-inspired Intelligence and a principal investigator at the Center for Excellence in Brain Science and Intelligence Technology (CEBSIT), where he does brain computational modeling.[33]
- ZHANG Tielin (张铁林). Zhang is an associate professor and principal investigator at CASIA, where he researches spiking neural networks, neuromorphic chips, memory, and cognitive decision-making. He served in lead roles in several national BI-AI projects.[34]
- ZHANG Zhaoxiang (张兆翔). Zhang is a professor at CASIA's Center for Research on Intelligent Perception and Computing and National Laboratory of Pattern Recognition, where he studies computer vision, pattern recognition, and "human-like learning."[35]

These same scientists surfaced in prior internet research on Chinese advanced AI.[36] Their appearance in peer-reviewed scientific publications is further support for China's claim to be working toward AGI.

3. ***The greatest concentration of Chinese AGI research is in the Beijing area.***

Most authors in the corpus claimed one or more affiliations.[37] The claims vary from paper to paper due to employment changes over the five-year period, journal policy, or name changes at the institution itself. The present study captured up to two affiliations per author per paper, three if foreign organizations were cited. In raw numbers, this equated to 2,943 claims that reduced to 370 unique Chinese and 62 foreign institutions. The institutions were entered into a spreadsheet with minor error correction.[38]

Analysis revealed that all but three of the 20 organizations highlighted in the July 2022 study as having an AGI focus, based on public statements, were represented in the corpus.[39] Whereas the mean number of papers for all Chinese institutions was 3.9, the average number of papers for these 17 organizations was 25.2. These statistics affirm

the validity of the earlier selections but also point to research at other Chinese institutions. The following are the 12 most prolific:

Table 1. Top Chinese Institutions Claimed by Authors of AGI-Related Publications

Organization	AGI publications
Peking University (PKU)	153
CAS Institute of Automation (CASIA)	152
Tsinghua University	139
University of the Chinese Academy of Sciences (UCAS)	104
Zhejiang University	61
Suzhou University	45
Nanjing University	41
Xi'an Jiaotong University	40
Beijing Institute for General Artificial Intelligence (BIGAI)	39
Dalian University of Technology	38
CAS Institute of Computing Technology	36
Beijing University of Technology	35

Source: CNKI database of Chinese academic journal articles

The outsized contributions of the CAS University, its Institute of Automation, and Peking and Tsinghua Universities take on special meaning in light of China's establishment in 2020 of a dedicated AGI institute—the Beijing Institute for General Artificial Intelligence (北京通用人工智能研究院), ninth-ranked in Table 1. BIGAI's operation is tightly integrated with AGI research at Tsinghua and Peking Universities, and it claims to have "close links" with other elements in the Beijing area, almost

certainly a reference to the Beijing Academy of Artificial Intelligence (北京智源人工智能研究院, BAAI).[40] The five institutions (see Figure 2) account for 587, or 70 percent, of the advanced (general) AI corpus.

Figure 2. Locations of China's Top Five AGI-Related Institutes within Beijing

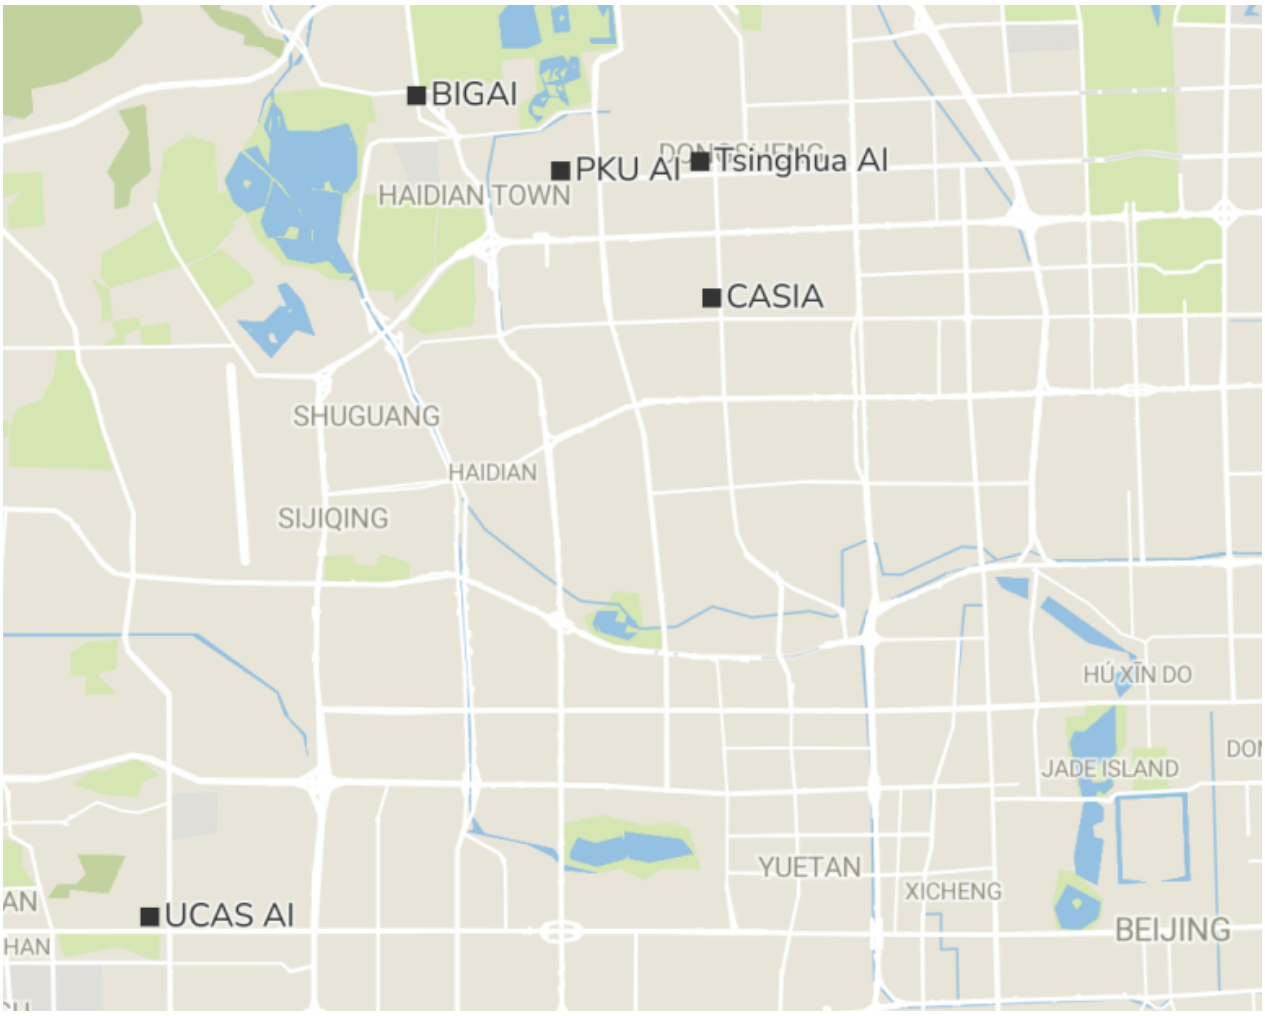

Source: Multiple (as described above)

Events support this study's finding that Beijing is—and will likely remain—China's main locus for AGI development or, at least, its most visible component. In October 2022, the Beijing Institute for General Artificial Intelligence was awarded national status as host to China's only "National Key Laboratory of General Artificial Intelligence" (通用人工智能全国重点实验室).[41] A year prior in November 2021, Beijing's municipal government announced the construction of a "Beijing General Artificial Intelligence Innovation Park" (北京通用人工智能创新园)—another first—in its Haidian District, home

to BIGAI, PKU, and Tsinghua University, which when completed in 2024 will serve as the platform (平台) for China's "AGI model."[42]

While acknowledging the Beijing-area concentration, given AGI's multi-disciplinary basis and the multiple paths through which it may be realized, the possibility of breakthroughs elsewhere in China cannot be ruled out. Limited data suggest that Beijing may be serving as China's AGI research hub for testing and deployment done elsewhere in China, in Wuhan especially.[43]

4. ***Global contributions support the research but are not its main drivers.***

Contributors to the present corpus of documents claimed affiliations with 10 foreign locations, shown in Table 2.[44]

Table 2. Numbers of Chinese AGI-Related Papers Citing Foreign Affiliations

Location	AGI publications with Chinese affiliation
Australia	11
Canada	7
Germany	1
Italy	1
Japan	5
Singapore	11
Taiwan	1
UAE	1
UK	18
United States	66

Source: CNKI database of Chinese academic journal articles

The high number of U.S. and UK affiliations is typical of Chinese AI research.[45] Also noteworthy are Australia's and Singapore's roles. Table 3 shows the number of institutions by country.

Table 3. Number of Foreign Institutions Listed by Authors of China AGI-Related Papers

Location	Number of institutions that contributed to AGI publications with Chinese affiliation
Australia	7
Canada	6
Germany	1
Italy	1
Japan	4
Singapore	4
Taiwan	1
UAE	1
UK	9
United States	25

Source: CNKI database of Chinese academic journal articles

The highest scoring foreign institutions were the University of Sydney in Australia (four claimed affiliations), National University of Singapore (six claims), Imperial College London (six claims), Carnegie Mellon and MIT (four claims each), and an astronomical 28 claims by scientists at UCLA, owing to BIGAI Director Zhu Songchun's (朱松纯) influence on UCLA's Center for Computer Vision, Cognition, Learning and Autonomous Robotics, which Zhu ran for 18 years (2002-2020).[46]

Although China's AGI research is supported by AI scientists and graduate students overseas, data show that the bulk of it is situated in Chinese institutions.

Caveats and Dependencies

This study's main goal is to assess claims by Chinese scientists about their AI research—work that leads by design or default to artificial general intelligence. While that goal has been met, the methodology used here may not fully capture the direction or scale of China's effort.

- Strategic research typically is not shared by nations or between all elements of a country's scientific community. So, although publications can point to a nation's scientific work, they may not deliver the full picture.[47]
- Furthermore, this study was unable to access full-text Chinese academic theses and dissertations.[48] Based on their titles, this genre potentially accounts for a higher ratio of on-topic studies than the journal publications.
- Resource constraints precluded a full review of the selected papers' bibliographies. Although sufficient material was collected to validate the study's thesis, exploiting this additional source would add to the number of titles captured.
- While the data supporting this paper's analytic judgments—a spreadsheet and corpus of 850 documents—are available for inspection, decisions affecting how this corpus was assembled—what papers met criteria for inclusion—rested on expert opinion that cannot be reproduced externally.
- China appears to be exploring multiple paths to AGI, including a potential approach not covered in this study, namely, cognitive sharing through BCIs.[49] Measuring China's investment in alternative approaches and how these efforts compare to global initiatives, while not a goal of this study, would be a useful exercise.

Finally, the path to AGI assumed in this paper—modeling human cognition—does not exhaust all possibilities and may not even be the quickest or most likely approach. The current debate among advocates of machine learning with their demonstrably successful large language models and proponents of cognitive or neuroscience-based approaches is far from resolved.[50]

Recommendations

1. A goal of this study is to demonstrate a protocol for alerting policymakers to threats and collaborative opportunities in AI and other technical areas.[51] While we accept the need for direct involvement in populating a foundational database, the labor-intensive approach used here is not sustainable and needs to be blended with automation and broader, institutional sponsorship for the project to run at scale.

2. The United States and its allies at present lack the means to gather information at the level of detail needed for rational decisions on how AI developments should be addressed. While we have no solution to AGI's *existential* challenges, what measures do emerge should be based on reliable information about nations' ability to field AGI in some form. The problem goes beyond China and AI to all breakthrough technologies. These developments must—and can be—monitored by a professional cadre using publicly available information ("open sources") in a defined federal-academic-policy institute construct. Urgent consideration should be given to its creation.

3. By the same token, we welcome contact with Chinese principals researching advanced AI. Most Chinese scientists active in this field have worked or will work abroad, in the United States especially, and many have dual affiliations.[52] These linkages should be captured, characterized, prioritized, and encouraged as a basis for outreach on shared concerns, including AI safety.

4. Similarly, both the United States and China need to accept the value of maintaining visibility into each other's scientific research. While recognizing a country's right to national defense, there are considerations that weigh heavily against an overabundance of secrecy in today's global research environment. At a minimum, there is a need for scientists everywhere to vet their work openly as a check against poor or fake scholarship—an issue that plagued China for decades—and *as quid pro quo for foreign collaboration*.

Hiding scientific research, for example, by restricting access to academic journals, may lead to false assumptions that devolve into a vicious cycle of measures and countermeasures—the U.S.-Soviet "missile gap" being the textbook example.[53] Pursuing this train of thought further, it is highly likely that one's inability to gauge the status and intent of a potential rival through open sources will lead to clandestine efforts to procure this same information—and more—driving science further underground to no one's long-term benefit.

5. Finally, we encourage U.S. policymakers contemplating AGI safeguards to bear in mind that this technology is sought by other nations aware of its strategic value and the importance of a "first-mover advantage." Although an unrestrained race to the top is risky, unilateral restrictions on AGI development,[54] trust-based agreements that cannot be verified, and one-sided adherence to ELSI/ELSA (ethics, legal and social implications/aspects) protocols are risky as well.[55]

Authors

William C. Hannas is CSET's lead analyst and formerly the Central Intelligence Agency's senior expert for China open-source analysis**.** Huey-Meei Chang is CSET's senior China S&T specialist. The two have co-authored and edited several books on Chinese technology policy, including *Chinese Power and Artificial Intelligence* (New York and London, Routledge, 2023). Max Riesenhuber, PhD, is professor of neuroscience at Georgetown University Medical Center and co-director of Georgetown's Center for Neuroengineering. Daniel Chou is a data scientist at CSET.

Acknowledgments

The authors are grateful to John Chen, Igor Mikolic-Torreira, and Jennifer Yunyan Wang for reading the manuscript and offering many helpful suggestions, and to Matt Mahoney for fact checking, Alex Friedland for copyediting and formatting, Ben Murphy for translation edits, and Tessa Baker for general oversight of the publication process.

Document Identifier: DOI: 10.51593/20220031

Endnotes

[1] Wm. C. Hannas, Huey-Meei Chang, Catherine Aiken, and Daniel Chou, "China AI-Brain Research" (CSET, September 23, 2020), 43, https://cset.georgetown.edu/publication/china-ai-brain-research/. "Brain-inspired" AI aims to improve the performance of AI systems by drawing on insights from neuroscience, whereas "cognitive AI" goes directly at the goal of AGI, that is to emulate human cognition, which can benefit from insights from neuroscience but is not constrained by them. LLMs are an example of this latter approach.

[2] Wm. C. Hannas, Huey-Meei Chang, Daniel Chou, and Brian Fleeger, "China's Advanced AI Research: Monitoring China's Paths to 'General' Artificial Intelligence," (CSET, July 2022), https://cset.georgetown.edu/publication/chinas-advanced-ai-research/ and Huey-Meei Chang and Wm. C. Hannas, "Spotlight on Beijing Institute of General Artificial Intelligence," (CSET, May 2023).

[3] PRC State Council, "New Generation AI Development Plan" (国务院关于印发《新一代人工智能发展规划》的通知), PRC State Council, 2017. China was pursuing brain-inspired AI research at least three years before its central government recognized AI as a standalone discipline. See Hannas, Chang, Aiken and Chou, "China AI-Brain Research."

[4] WU Guozheng, XIAO Bin, ZHAO Ruizhen, CHEN Ting (吴国政，肖斌，赵瑞珍，陈厅), "The NSFC funding situation and disciplinary development prospects of artificial intelligence in 2021" (2021 年度 NSFC 人工智能学科基金项目申请资助情况及学科发展展望), *CAAI Transactions on Intelligent Systems* (智能系统学报) 16, no. 6 (2021): 1166-1171.

[5] Wm. C. Hannas and Huey-Meei Chang, "China's STI Operations," (CSET, January 2021), https://cset.georgetown.edu/publication/chinas-sti-operations/. "OSINT" is an acronym for "open source intelligence," i.e., analysis based on publicly available information (PAI).

[6] "The U.S. Government recognizes that open-source data is an underutilized resource in advancing science and technology (S&T) dominance… As U.S. adversaries march to the forefront of S&T competition, their advancements are visible through their contributions to S&T literature.… This exposure leaves much of adversarial S&T development activity abroad open to discovery through Open Source Intelligence (OSINT) techniques developed in academia, industry, and the private sector." Loren Blinde, "DIU seeks solutions for Blue OSINT," *Intelligence Community News*, January 9, 2023, https://intelligencecommunitynews.com/diu-seeks-solutions-for-blue-osint/.

[7] Note China's establishment in March 2023 of a new state program called "AI for Science" (人工智能驱动的科学研究, literally "AI-driven scientific research") to support discovery in key fields such as biology, pharmacology, and new materials. See "Ministry of Science and Technology launches special program for 'AI-driven scientific research'" (科技部启动"人工智能驱动的科学研究"专项部署工作), *Huanqiu.com* (环球网), March 28, 2023, https://m.huanqiu.com/article/4CGlNqAw2KO.

[8] E.g., authors, laboratories, funding codes, citations etc. The initial search terms typically miss relevant work. A second, metadata-based search can double the document count.

[9] Data on Chinese brain-computer interfaces (BCI) will be treated in a separate paper (see below).

[10]中国知网, a private Chinese S&T document aggregator, https://www.eastview.com/resources/cnki-faq/. Searches were done through the authors' Georgetown University Library subscribed service.

[11] "Wanfangdata" OR "万方数据"; "CNKI" OR "中国知网"; "CQVIP" OR "维普".

[12] The 2018-2022 timeframe is intended to capture scholarship that appeared after the State Council's 2017 "New Generation AI Development Plan." A few papers published in early 2023 were also retrieved.

[13] As former Chinese leader Deng Xiaoping said, "It doesn't matter whether a cat is black or white, as long as it catches mice." (不管黑猫白猫，能捉到老鼠就是好猫).

[14] Hannas and Chang, et al., "China's Non-therapeutic BCI Research: Alternate Paths to Cognitive Augmentation and Control," CSET, (forthcoming, 2023). Some 4,078 Chinese-authored papers on BCI were identified, only a few of which had a specific AGI focus.

[15] The first of the five parameters (SME assessment) was given a weight of 2; the other four carried a weight of 1 each. Journal rankings (2) using two measures were provided by CNKI—the present study used their average. Papers funded by the National Natural Science Foundation of China, the National Social Science Fund of China, or a national level ministry fulfilled category (3). Citation frequency (4) and number of downloads (5) were adjusted for year of publication. Categories (2), (4), and (5) were fulfilled if the paper ranked in the top half of the distribution. Final selection was based on a paper earning 3 or more points, e.g., identified by the study's SME as significant (2 points) + high number of downloads (1 point), or highly cited (1 point), top level journal (1 point) + NNSF-funded (1) point, and so on. The exercise was a time-saving measure that took into account limited resources. Ideally the (co)authors of *all papers* would be searched for additional work that eluded capture in earlier runs.

[16] *Artificial Intelligence* (人工智能) is published by the China Center for Information Industry Development (CCID; 中国电子信息产业发展研究院；赛迪).

[17] Chinese names have two or three syllables, each syllable represented by a written character. In Chinese, a character (字) has one (rarely more) sound(s) associated with it, but a given syllable can be expressed by multiple characters depending on its meaning. Extreme examples are the syllables *yi* and *shi*, each of which are associated with many dozens of characters. Hence when a name is written in characters—the norm in China—the information is one or more *orders of magnitude* greater than when the author's name is alphabetized in China's official *pinyin* notation, used in non-Chinese journals. Complicating matters further, Chinese Mandarin manages with just 1,280 unique syllables, including phonemic tone, which typically is *not* rendered in the *pinyin* transcription. So the set of available symbols

(written in *pinyin*) is a third of that number, compared to some 7,000 commonly used characters. Worse yet, characters used as surnames draw from a pool of about 200. This is balanced by fewer restrictions on the selection of characters for given names that make up a full name's second (and optionally third) element(s). So the information needed for discrimination is mostly carried in the given name, not the surname, which precedes the given name in order of appearance. Many English language journals, however, show only the author's surname followed by initial(s), which can obscure the identity of the contributor even when affiliations are provided. For example, 朱松纯, whose character name is unambiguous, becomes "Zhu Songchun" or "Zhu, S.C." or even "Zhu, S." in a paper's author byline. There are ways around the problem but they are convoluted, time-consuming, and can fail for less well-known authors especially. In the present study, 30 percent of the *pinyin* names of first authors in English language fora could not be matched with characters by this study's authors, one of whom wrote a definitive book on the problem, and who have been dealing with the issue professionally for decades. William C. Hannas, *Asia's Orthographic Dilemma* (Honolulu: University of Hawaii Press, 1997).

[18] The date range reflects the availability of local resources.

[19] Formerly the Journal of Zhejiang University Science C (Computers & Electronics) in English.

[20] *International Journal of Automation and Computing* (国际自动化与计算杂志) and *Machine Intelligence Research* (机器智能研究, after 2022), both published in English by CASIA.

[21] Hannas, Chang, Chou, and Fleeger, "China's Advanced AI Research."

[22] Hannas, Chang, Aiken, and Chou, "China AI-Brain Research."

[23] Kai-Fu Lee, *AI Superpowers* (New York, Houghton Mifflin Harcourt, 2018).

[24] PRC State Council, "New Generation AI Development Plan" (国务院关于印发《新一代人工智能发展规划》的通知), PRC State Council, 2017.

[25] Author information is derived from the authors' bios, academic CVs, and other open source material and is current as of May 2023.

[26] Gao Wen's CV, https://ieeexplore.ieee.org/author/37275735100; https://cfcs.pku.edu.cn/people/directors/wengao/index.htm.

[27] Huang Tiejun's CV, https://www.ai.pku.edu.cn/info/1139/1243.htm; https://2022.baai.ac.cn/about.

[28] Liu Jianwei's CV, https://saiconference.com/intellisys2023/CommitteeProfile/8ee3d253-2680-4ab1-9355-ee80c1b13031; https://www.cup.edu.cn/cise/szdw/fjs1/170307.htm.

[29] National University of Defense Technology (国防科技大学) in Changsha, administered by the Ministries of National Defense and Education. See Wang Feiyue's CV at http://www.impcia.net/expert/details_69.html; https://ieeexplore.ieee.org/author/37277656000; https://www.kmust.edu.cn/info/1011/4362.htm.

[30] Wu Fei's CV, https://person.zju.edu.cn/en/wufei; https://person.zju.edu.cn/wufei.

[31] XU Bo's CV, https://people.ucas.ac.cn/~xubo.

[32] Yang Yaodong's CV, https://www.yangyaodong.com/.

[33] Zeng Yi's CV, https://people.ucas.ac.cn/~yizeng. CEBSIT is an umbrella organization for 39 Chinese research institutes. CAS's Institute of Neuroscience (神经科学研究所) and CASIA are its main supporting units.

[34] Zhang Tielin's CV, https://people.ucas.ac.cn/~tielin.zhang; https://bii.ia.ac.cn/~tielin.zhang/.

[35] Zhang Zhaoxiang's CV, https://zhaoxiangzhang.net/.

[36] See Hannas, Chang, Aiken, and Chou, "China AI-Brain Research;" Hannas, Chang, Chou, and Fleeger, "China's Advanced AI Research;" and Wm. C. Hannas and Huey-Meei Chang, eds., *Chinese Power and Artificial Intelligence* (New York and London: Routledge, 2023).

[37] A few papers inexplicably lacked author affiliations, although their authors' units were shown in other papers. The omissions were likely oversights.

[38] Regularizing institutional nomenclature is a mandatory step for displaying names usefully in a database.

[39] The State Key Laboratory of Brain & Cognitive Science's (脑与认知科学国家重点实验室) and Shanghai Center for Brain Science and Brain-inspired Intelligence's (上海脑科学与类脑研究中心) absence from the corpus reflects a decision to exclude papers where AI played only a supporting role. The other outlier was Horizon Robotics (地平线机器人), a smaller operation with offices in Beijing and Silicon Valley.

[40] "Toward General Artificial Intelligence" (迈向通用人工智能), 44, https://www.bigai.ai/flip-book/. BIGAI issued this 50-page pamphlet in late 2022. BAAI is led by Huang Tiejun, a strong advocate of artificial super intelligence (ASI), who is also vice-dean of PKU's Institute for Artificial Intelligence (北京大学人工智能研究院)—an integral part of the BIGAI consortium. Chang and Hannas, "Spotlight on Beijing Institute of General Artificial Intelligence."

[41] "Brief Introduction to Peking University's School of Intelligence Science and Technology" (北京大学智能学院简介)," Beijing Institute for General Artificial Intelligence, January 9, 2023,

https://www.bigai.ai/news/school-of-intelligence-science-and-technology/. See Chang and Hannas, "Spotlight on Beijing Institute of General Artificial Intelligence" for more details on the key lab's attribution and nomenclature.

[42] "Beijing General Artificial Intelligence Innovation Park starts construction and is expected to be completed in November 2024" (北京通用人工智能创新园开建预计 2024 年 11 月竣工), *Haidian News* (海淀报), November 26, 2021, http://zyk.bjhd.gov.cn/ywdt/rdgz/202111/t20211129_4497925.shtml?type=computer.

[43] Wuhan's evolution as an AGI center is described in Hannas, Chang, Chou, and Fleeger, "China's Advanced AI Research," 32-35. Additionally, in late 2022, a PKU Wuhan Institute of Artificial Intelligence was established as a joint venture of Peking University and the Wuhan city government. Zhu Songchun, BIGAI's director, serves as its chief scientist, leading a team that oversees regional AGI deployment. "Peking University Wuhan Institute of Artificial Intelligence Recruitment Spring 2023" (北大武汉人工智能研究院 2023 年春季招聘), Student Career Center of PKU, March 7, 2023, https://scc.pku.edu.cn/employment_22e9034586a29d2d0186baf675fa16b1_0.html; "Peking University Wuhan Institute of Artificial Intelligence Settles into Optics Valley, Zhu Songchun serves as chief scientist" (北京大学武汉人工智能研究院落户光谷，朱松纯任首席科学家), Zhongguo Guanggu (中国光谷), December 22, 2022, http://news.cjn.cn/cqpd/jzgg_19988/202212/t4392646.htm.

[44] Papers citing author affiliations to China-based multinational corporations (NVIDIA, Intel, Microsoft Research Asia—13 papers totally) are an intermediate category not reflected in the tables.

[45] "Overall, U.S.-China collaborations on AI research have quintupled since 2010 and totaled 9,660 papers in 2021—much faster than the increase in collaborations between any other two nations. Collaborations between the United States and United Kingdom, the second most prolific source of cross-border research, increased almost threefold to 3,560 papers." Edmund L. Andrews, "China and the United States: Unlikely Partners in AI," *Stanford University Human-Centered Artificial Intelligence*, March 16, 2022. https://hai.stanford.edu/news/china-and-united-states-unlikely-partners-ai.

[46] Chang and Hannas, "Spotlight on Beijing Institute of General Artificial Intelligence."

[47] China itself operates the world's most comprehensive and advanced open-source STI (science and technology intelligence) operation and acknowledges OSINT as its most productive source. See Hannas and Chang, "China's STI Operations: Monitoring Foreign Science and Technology through Open Sources," GU/CSET, January 2021.

[48] The Georgetown University library provides access to the full-texts of Chinese academic journal titles but does not subscribe to Chinese theses and dissertations.

[49] Hannas, Chang, Chou, and Fleeger, "China's Advanced AI Research."

[50] Pengcheng Zhou (周鹏程), "A twitter debate: whether artificial intelligence needs neuroscience" (一场 twitter 争论：人工智能是否需要神经科学), Zhihu.com (知乎网), October 24, 2022, https://zhuanlan.zhihu.com/p/576570463; Gary Marcus, "Deep Learning Alone Isn't Getting Us to Human-like AI," *Noema*, August 11, 2022, https://www.noemamag.com/deep-learning-alone-isnt-getting-us-to-human-like-ai/; Gary Marcus, "Deep Learning Is Hitting a Wall," *Nautilus*, March 10, 2022, https://nautil.us/deep-learning-is-hitting-a-wall-238440/; Melissa Heikkilä and Will Douglas Heaven, "Yann LeCun Has a New Vision for the Future of AI," *MIT Technology Review*, June 24, 2022, https://www.technologyreview.com/2022/06/24/1054817/yann-lecun-bold-new-vision-future-ai-deep-learning-meta/.

[51] See Hannas, Chang, Chou, and Fleeger, "China's Advanced AI Research" for the details of a proposed AGI watchboard.

[52] CSET survey conducted through Qualtrics between January 5 and March 4, 2020. Hannas, Chang, Aiken, and Chou, "China AI-Brain Research," 42.

[53] https://www.cia.gov/readingroom/collection/what-was-missile-gap.

[54] As a case in point, at the same time AI experts, predominantly in the United States and Europe, went on record supporting a moratorium on AI research out of concern for AGI's potential hazards, Chinese President Xi Jinping instructed his Politburo to "attach importance to the development of AGI" (要重视通用人工智能发展). "The Political Bureau of the CPC Central Committee held a meeting to analyze and study the current economic situation and economic work. Xi Jinping, general secretary of the CPC Central Committee, presided over the meeting" (中共中央政治局召开会议 分析研究当前经济形势和经济工作 中共中央总书记习近平主持会议), *Xinhua*, April 28, 2023, http://www.gov.cn/yaowen/2023-04/28/content_5753652.htm.

[55] ELSI (in the United States, ELSA in Europe) is an acronym for research on "non-technical issues that arise when developing emerging science and technologies and implementing them in society." See Atsuo Kishimoto, "What Is ELSI?" https://elsi.osaka-u.ac.jp/en/what-is-elsi.