

Issue Brief

Spotlight on Beijing Institute for General Artificial Intelligence

China's State-Backed Program
for General Purpose AI

Authors

Huey-Meei Chang

William C. Hannas

Executive Summary

China's Beijing Institute for General Artificial Intelligence (BIGAI), established with state backing in 2020, aims openly at artificial general intelligence (AGI) and is assembling the talent and organizational means to pursue it.

The project's core is an elite team of Chinese- and U.S.-educated scientists managed by former University of California, Los Angeles (UCLA) researcher Zhu Songchun, whose work in precursor disciplines, professional network, and openness to methodological alternatives lend credibility to the project.

The present study—an introduction to BIGAI's goals, staffing, and research—situates this AGI project in the context of China's overall work toward advanced artificial intelligence.

- BIGAI by choice is not pursuing the massively large natural language models championed by OpenAI, Google, and other U.S. and British companies.
- BIGAI focuses instead on developing AGI through alternative “small data, big task” research on brain cognition and neuroscience.
- The organization has recruited some 30 top scientists educated at leading U.S. and UK research universities, several of whom were trained under U.S. government programs.
- Zhu described his AGI project to high-level state bodies as being on a par with China's storied “two bombs, one satellite” programs in terms of its importance.

Given the strategic impact of a successful Chinese AGI program, this study encourages U.S. and allied government policymakers to pay closer attention to China's AI development, through open sources especially, as a foundation for practical engagement.

Introduction

In September 2020, Chinese newspapers announced the establishment in Beijing of a large-scale institute dedicated to creating artificial general intelligence (AGI).¹ This has long been regarded in AI circles as a holy grail—or an unobtainable fantasy, or humankind's final invention.²

The initiative was not unprecedented. Efforts to imbue AI with general capabilities modeled loosely or explicitly on human cognition began years ago, in China and elsewhere.³ Current developers include OpenAI in the United States, DeepMind in the UK, to cite two prominent examples, and an assortment of Chinese companies and labs operating largely under the radar.⁴

What distinguishes this new Chinese project—beyond size, leadership, and state backing—is its open embrace of “general purpose” artificial intelligence as its founding mission, enshrined in the name Beijing Institute for General Artificial Intelligence (北京通用人工智能研究院). Recent materials published by the institute and statements from its major figures allow a closer look at this unique organization.

Overview of the institute

BIGAI began operations in October 2020 under the sponsorship of China's Ministry of Science and Technology, its Ministry of Education, and Beijing's municipal government. The location in Beijing's Haidian District, adjacent to Peking and Tsinghua Universities, facilitates interaction between the three institutions, whose programs are aligned.⁵ Claims that the facility will be staffed by 1,000 researchers “drawn from China and abroad” are consistent with the building's size and reinforced by a literal reading of its acronym BIGAI.⁶



Source: Beijing Institute for General Artificial Intelligence.⁷

Beyond its own hires, BIGAI draws on AI researchers from China's two top-ranked academic institutions, Tsinghua University and Peking University (PKU). The latter encompasses the PKU Institute for Artificial Intelligence (北京大学人工智能研究院) and its subordinate Institute for General Artificial Intelligence (通用人工智能研究所), as well as PKU's School of Intelligence Science and Technology (北京大学智能学院), which includes the Wangxuan Institute (王选所).⁸ Tsinghua University is affiliated with BIGAI through its Department of Automation (清华大学自动化系) and a planned Institute for General Artificial Intelligence (通用人工智能研究院).⁹

The BIGAI complex is supported by an undergraduate "AGI experimental class" (通用人工智能实验班) populated by elite students in Tsinghua's Department of Automation and PKU's Yuan Pei College (元培学院).¹⁰ In addition, BIGAI has collaborative links with other institutions in the Beijing area.¹¹

BIGAI's status as China's premier state-sponsored AGI research organization is underscored by the creation of a "National Key Laboratory of General Artificial Intelligence" (通用人工智能全国实验室) within BIGAI, the only such state key laboratory to bear the "AGI" name.¹²

Goals and research

“Artificial general intelligence” (AGI)—the term and the construct—acquired pariah status outside China as a result of early disappointments and unsupported claims by pundits lacking appreciation of its difficulty. Global experts cautiously estimated that AGI’s arrival would be decades away, with some even doubting it in principle.¹³ For scientists already struggling with “AI funding winters,” the term became a symbol of naivete and a label to be avoided. Even today researchers use phrases such as “general purpose artificial intelligence” to preempt criticism from skeptical peers.¹⁴

China’s history with the concept radically differs. Due to its language and isolation from the global politics of AI research,¹⁵ China never struggled with an AGI identity crisis. The Chinese term for AGI 通用人工智能 (tōngyòng réngōng zhìnéng), in use since Goertzel and Pei Wang reintroduced “AGI” in or about 2002,¹⁶ translates to “general purpose artificial intelligence,” identical to today’s preferred expression, releasing China from the stigma attached to the English acronym. Despite this natural immunity, Chinese scientists prefer “AGI” in their discourse and have no qualms publicly relating to it.¹⁷

Nowhere is this fact clearer than at BIGAI. The institute is not shy about its aim to replicate all aspects of human cognition, in the AGI concept’s original sense.¹⁸ This goal is emphasized in BIGAI’s use of the character 通 (“general”) decoratively on the building’s exterior, and in a forced interpretation of the character’s parts used in promotional literature to encourage Chinese identification with AGI and ownership of the mission.¹⁹

At the Beijing Institute, AGI is an all-encompassing principle, ideologically and as reflected in its slate of tasks. These include building new AI theories and paradigms, resolving “cross-media bottlenecks” (跨媒体瓶颈), and creating a general purpose (通用) artificial intelligence operating system, general purpose agents, and associated training and testing platforms.²⁰

These “platforms” (平台) support three-dimensional simulation training, data acquisition, autonomous robot hardware development, high-performance graphics processing unit computing, cross-modal awareness, audio-visual psychological experiments, and “cognitive mechanism public experiments” (认知机理公共实验),²¹ as well as research in the following areas:

- Vision
- Text reasoning
- Syntax and semantics
- Information extraction
- Common sense understanding
- Text generation and intelligent dialogue
- Multimodal fusion embodied intelligence

Additional research addresses related topics in machine learning and cognitive science, aimed at “laying the foundation for building general-purpose agents that interact and collaborate naturally with humans.”²²

These agents will “have a wide range of human-like cognitive abilities, common sense learning, and reasoning capabilities” meant to “approximate human thinking” (更接近人脑思维).²³ They will perform “hierarchical reinforcement learning” based on a multi-agent approach that involves “communication, cooperation, and competition” as well as understanding “social common sense and norms” (社会常识与规范), and that will have emotional intelligence to “better integrate into complex social environments.”²⁴

PKU’s Institute for Artificial Intelligence, which, as noted, is allied with BIGAI, is positioned to support AGI development, as evidenced in the following research goals expressed as task “layers” (层):

- Key domain layer: develop computer vision, natural language processing, computational cognition, common sense reasoning, multi-agent bodies, robotics, and machine learning;
- General platform layer: build a general artificial intelligence system platform and a large task test platform;
- Interdisciplinary support layer: create intelligent system software and intelligent brain-like chips; develop visual perception, foundational mathematical research, artificial intelligence governance, computational social science, and intelligent medical care.²⁵

The PKU institute, among other task-specific laboratories, also has a Cognitive Reasoning Lab that researches abstract and interactive reasoning, analogy, causality, and number sense.²⁶

In practical terms, the composite enterprise will “provide physical embodiment for general artificial intelligence” aimed at “self-sufficiency in open environments” and at exercising “initiative and adaptability with a minimum of human intervention,” so as to “perform complex, dynamic and diverse real life tasks.”²⁷ BIGAI’s metaphysical goal of a “unified theory” (统一理论) of AI and cognition is no less ambitious. These aims will be realized by shifting from a data-driven to a task-driven paradigm, a key element of director Zhu Songchun’s (朱松纯) philosophy.²⁸

Staffing and “talent” acquisition

BIGAI draws its talent from open recruitment of domestic and foreign experts, from the faculty and students of its affiliated departments at Peking and Tsinghua Universities, and from intensive in-house training of select undergraduates, attracted by promises of leading the new future of AI and “going where no one has gone before” (无人之境).²⁹ One notice solicits:³⁰

- **Researchers** committed to studying next generation AI systems in a wide range of fields such as computer vision, machine learning, cognitive computing, common sense reasoning, robotics, multi-agents and large-scale task simulation.
- **Engineers** for the development of computer vision, natural language processing, machine learning, cognitive computing, common sense reasoning, robotics, multi-agent and large task simulation.
- **Visiting scholars** including professors, scholars, and researchers in the field of artificial intelligence worldwide.
- **Postdoctoral researchers** for a two-year (renewable) joint training program in cooperation with Tsinghua University and Peking University.

BIGAI also staffs a Cooperation Center to “create a good environment for external cooperation” and “recruit outstanding talent for related fields,” and it manages internships for doctoral, master's, and undergraduate degree students in top universities at home and abroad.³¹

The biographies below illustrate the talent BIGAI brings to its mission (a discussion of BIGAI’s director, Zhu Songchun, follows):

- FAN Lifeng (范丽凤). A former statistician under Zhu Songchun at the University of California, Los Angeles (UCLA), Fan followed Zhu back to China to establish BIGAI. Fan apparently took a pay cut (“the enterprise research institutes pay better”), but states that she prefers BIGAI’s work environment and academic freedom.³²
- HUANG Siyuan (黄思远). Huang interned at DeepMind and Facebook’s Reality Lab. He is now a team lead at BIGAI, where he studies computer vision, machine learning, cognition, and robotics. Huang’s goal is to develop “generalizable models for general-purpose agents.”³³
- LIANG Yitao (梁一韬). Liang, who holds a Ph.D. from UCLA, returned to China to join BIGAI’s “small data, big tasks” project, which he calls “rare in this world.” Professor Liang’s impression about returning to China is that the quality of students there is “higher than students in North America’s famous schools (高于北美名校学生).”³⁴
- LIU Hangxin (刘航欣). Liu trained in the United States under grants from the Defense Advanced Research Projects Agency (DARPA) and the U.S. Office of Naval Research (ONR). He believes “there are fundamental representations or cognitive architectures underpinning the intelligent behavior of humans.” His aim is to reproduce them in robots.³⁵
- LIU Tengyu (刘腾宇). Liu, who earned a master’s and Ph.D. under Zhu Songchun’s mentorship at UCLA, now works in BIGAI’s General Vision Lab. His “goal is to create intelligent agents that can interact with virtual or physical environments just like humans.”³⁶
- PENG Yujia (彭玉佳). Peng, a UCLA alumna, researches the “intersection of psychology and artificial intelligence.” Her impression upon returning to China is that AGI there is taken more seriously (“seen as important”). Peng’s research covers a spectrum of cognitive processes.³⁷
- YANG Yaodong (杨耀东). Yang researches reinforcement learning and multi-agent systems at PKU’s Institute for Artificial Intelligence. Before joining BIGAI, he was a principal research scientist at Huawei, UK, and assistant professor at King’s College London.³⁸

- ZHANG Muhan (张牧涵). Zhang is an assistant professor at PKU and a researcher at BIGAI. He holds a Ph.D. from Washington University and was a research scientist at Facebook, where he worked on deep learning. He now researches machine learning/DL and graph neural networks.³⁹
- ZHENG Zilong (郑子隆). Zheng, a UCLA alumnus who has returned to China, studies the “intersection of statistical machine learning, vision-language modeling, and cognition.” His areas of specialization are generative modeling, multimodal understanding, and cognitive language reasoning.⁴⁰
- ZHU Yixin (朱毅鑫). Zhu, who holds a Ph.D. from UCLA, now directs PKU’s Cognitive Reasoning Lab, where he studies abstract, visually grounded, and interactive reasoning. Zhu hopes to integrate “high-level common sense (functionality, affordance, causality, intent) with visual and haptic (touch) perception.”⁴¹

BIGAI’s recruitment efforts, supported by institutional reforms and innovations provided by the Beijing city government, are meant to endow “strategic scientists” with personal, financial, and academic autonomy.⁴² These perks and accommodations are cited as factors for relocating to China by these and other BIGAI researchers.

The “Tong Ban” (通班) class of new AGI professionals

BIGAI’s third source of talent is an “AGI experimental class” (通用人工智能实验班) of elite undergraduates from Peking and Tsinghua Universities.

The institute’s leadership reasons that to build AGI, it is not enough to hire skilled domestic or foreign talent. Rather, “fresh blood” (新鲜血液) needs to be inculcated in the goals and needs of the enterprise, specifically:

The ultimate concept of general artificial intelligence is to endow machines with a series of comprehensive capabilities such as perception, cognition, decision-making, behavior, and execution, which requires cultivating professionals with interdisciplinary knowledge.⁴³

PKU’s Yuan Pei College and Tsinghua’s Institute of Automation emphasize interdisciplinary research. Beyond the students’ elite academic status, another draw for BIGAI is that they have had less exposure to mainstream values and establishment beliefs. BIGAI scientists educate the undergraduates in the program’s needs and goals.

Following is a list of several training courses offered to BIGAI students in spring 2022:⁴⁴

- A unified framework for abstract learning
- Cognitive state understanding based on body language
- Cooperative multi-agent reinforcement learning
- Exploration of physical properties of objects
- Joint vision-language understanding
- Mobile robot scene reconstruction and navigation
- Multi-agent task planning based on theory of mind
- Multimodal neural-symbolic question answering system
- Multi-object selection based on reinforcement learning
- Question-answering systems in embodied environments
- Real-time prediction of human action intention
- Reinforcement learning based on physical rules
- Research on real-time multi-object state fusion
- Scene and task generation for embodied intelligence
- Unsupervised semantic role labeling for language
- Visual logic grounding

These courses acclimate China's brightest AI students to precursor fields that the institute believes are needed to achieve AGI. They indicate the types of research carried out at the BIGAI complex. The topics address areas generally viewed as AGI "bottlenecks."⁴⁵



Source: PKU and Tsinghua “AGI Class” (通班).⁴⁶ The smaller characters read “Build an ace-in-the-hole-technology army, create general intelligence.”

Zhu Songchun: background and philosophy

BIGAI director Zhu Songchun (朱松纯) returned to China in September 2020 after 28 years in the United States to found the Beijing AGI institute with national and municipal support. His incentive package included appointments as dean of PKU's Institute of Artificial Intelligence, dean of its School of Intelligence Science and Technology,⁴⁷ and head of Tsinghua's new AGI Institute.⁴⁸ These posts ensure that Zhu's vision of AGI will dominate the enterprise.

A review of Zhu's background adds perspective to his choices. Zhu graduated from Hefei's University of Science and Technology of China (中国科学技术大学) in 1991 with a degree in computer science. The provincial university and city itself have played outsized roles in China's high-tech civilian and military development. Zhu went on to earn a doctorate from Harvard in 1996 and do stints at Brown, Stanford, and Ohio State before being appointed to UCLA in 2002. Four years later he was a professor, serving as the head of UCLA's Center for Vision, Cognition, Learning and Autonomous Robotics.⁴⁹

Zhu's specialties are computer vision and statistical modeling. An Institute of Electrical and Electronics Engineers (IEEE) fellow, he has written more than 350 papers in his two main fields and in cognitive science, language understanding, robot autonomy, and common sense reasoning.⁵⁰ Zhu's contributions to computer vision earned him high-level honors, including early selection to China's Thousand Talents Program.⁵¹

His contrarian advocacy of "small data, big task" (小数据, 大任务) was borne of impatience with the mainstream big data approaches that Zhu at one time professed. To illustrate his former beliefs (and frustrations), following is the slide Zhu presented at an October 2019 conference at the Beijing Academy of Artificial Intelligence, 11 months before exiting UCLA.⁵² The banner reads: "The World's Earliest Big Data Labeling Team: Hubei Lotus Hill Institute (LHI) 2005."



Source: The World's Earliest Big Data Labeling Team: Hubei Lotus Hill Institute (LHI) 2005.⁵³

Zhu's point, reiterated in many presentations, was that he once sought big data solutions for AI, in his field of computer vision especially, but abandoned them as unworkable.⁵⁴ By 2009, Zhu's disenchantment with AI's non-interpretability and inability to generalize from specific to broad tasks led him to "switch to cognitive science" for solutions to AI's conundrums.⁵⁵ His transition from the mainstream approach to AI in the United States to a brain-inspired model also followed from his belief that AI's end goal was, and is, *human-like* AGI. In his view, AGI thus understood "was the original intention and is the ultimate goal of intelligence science; it is the commanding height of S&T [science and technology] and the only way to go."⁵⁶

Zhu's views on the necessity and inevitability of AGI are expressed in multiple fora. The following is from a talk on safe AI hosted by the Association for the Advancement of Artificial Intelligence on February 13, 2023, that featured Zhu and UC Berkeley's Stuart Russell:

When I formed the Beijing Institute for General Artificial Intelligence, a new research and development institution in China, I specifically decided to use AGI as part of the institution's name to distinguish it from dedicated AI. . . . The goal is to achieve a general intelligent body with autonomous perception, cognition, decision-making, learning, execution, and social collaboration capabilities, in line with human emotions, ethics, and moral concepts.⁵⁷

Zhu went on to describe how humans and AGI would "align" on four levels, by sharing:

- Representation, including a common understanding of the world
- Core knowledge, such as a common sense of physics, causal chains, and logic
- Social norms, namely, the norms of human society and appropriate behavior
- Values, i.e., alignment with human moral principles

At the same talk, he reiterated an argument prevalent among China’s AGI advocates:

Modern scientific research shows that life on Earth is constantly evolving, from inorganic matter to organic matter, single-celled, multicellular organisms, then to plants, animals, and finally to intelligent creatures like humans. This reveals that from “physical” to “intelligent” is a continuous evolutionary process. We have no reason to think humans will be the end of this evolutionary spectrum. This also indicates that it is possible for general intelligence to surpass humans in the future.⁵⁸

On March 4, 2023, Zhu, a member of the Chinese People’s Political Consultative Conference, addressed that body and China’s National People’s Congress with an appeal to fast-track China’s development of general purpose AI. Zhu likened AGI to China’s “two bombs and one satellite” (两弹一星), a phrase used to describe a Chinese project of paramount importance and national pride.⁵⁹ In Zhu’s words:

If China can be the first to achieve a truly generally intelligent entity, it will become the “ace in the hole” of international scientific and technological competition.⁶⁰

He went on to describe his vision of what is needed on a national scale:

Coordinated and deployed by the central government, gather the advantages of talent, scientific research, industry and capital, standardize the development path, and issue a roadmap for the development of general artificial intelligence in China as soon as possible. Establish testing standards, and plan the underlying cognitive architecture, core algorithms, operating systems, and programming language, architecture and chips, forming an innovation chain of “university, research, production and application,” and seize the commanding heights of global technology and industrial development in the emerging strategic field of general artificial intelligence.⁶¹

Zhu’s rejection of the big data (large language model) approach to advanced AI is shared by other Chinese scientists, such as Chinese Academy of Sciences (CAS) academician Zheng Nanning (郑南宁), chairman of the Chinese Society of Automation (中国自动化学会). He stated: “The way to solve the current problems of artificial

intelligence is to return to the fields of brain cognition and neuroscience, and obtain inspiration and enlightenment from the human brain and wisdom.”⁶²

Zhu’s passion for AGI is also taken up by other prominent Chinese researchers, many of whom share his brain-inspired (类脑) approach and belief that AGI will not stop at the human level.⁶³

Foreign connections and dependencies

Zhu’s focus on building a unique Chinese identity for AGI, and his non-mainstream outlook on how to achieve it, belie the fact that the wellspring for his enterprise was—and to an unknown extent still is—the United States and Britain.

Zhu spent the last two-thirds (2002–20) of his 28 years abroad directing a team of researchers at UCLA. During that time, the lab received multiple U.S. funding awards from DARPA and ONR through its Multidisciplinary University Research Initiative.⁶⁴

Thus Zhu’s 2020 return to China, while problematic for U.S. national security, also entails a long-term research integrity problem.⁶⁵ A recent source puts the number of foreign-trained scientists recruited to BIGAI at “over 30,”⁶⁶ many of whom were members of Zhu’s American-trained UCLA team, while those still at UCLA continue to coauthor with staff at BIGAI.⁶⁷

For example, six of the seven authors of a 2020 DARPA- and ONR-funded paper are now at BIGAI.⁶⁸ A paper published in 2022 with DARPA funding was coauthored jointly by UCLA and BIGAI scientists, three of whom acknowledged competing interests.⁶⁹ Another 2020 DARPA- and ONR-funded paper had coauthors from UCLA, BIGAI, and Massachusetts Institute of Technology (MIT).⁷⁰ Following are titles of other U.S. government-funded studies that were coauthored by returnees now at BIGAI:

- “Interactive Robot Knowledge Patching Using Augmented Reality,” 2018 IEEE International Conference on Robotics and Automation (ICRA).
- “Interpretable Convolutional Neural Networks,” 2018 IEEE Conference on Computer Vision and Pattern Recognition (CVPR).
- “Learning Perceptual Inference by Contrasting,” NeurIPS 2019 Spotlight.
- “LEMMA: A Multi-view Dataset for Learning Multi-agent Multi-task Activities,” European Conference on Computer Vision (ECCV), 2020.

- “Patching Interpretable and-or-Graph Knowledge Representation Using Augmented Reality,” *Applied AI Letters* 2, no. 3 (October 2021).

ONR’s Multidisciplinary University Research Initiative, which Zhu headed on two occasions, has representation from Carnegie Mellon, MIT, Stanford, UCLA, Yale, and the University of Illinois on the U.S. side; and Birmingham, Glasgow, Leeds, and Oxford Universities on the UK side.⁷¹ British links are also evidenced in a 2022 NeurIPS paper, five of whose authors hail from UK institutions and the remaining three, including one who worked at King’s College London, from BIGAI and other Chinese AI institutions.⁷²

Whose AGI—and does it matter?

While recognizing BIGAI’s foreign antecedents and links to international scholarship, it is easy to overlook the fact that Zhu’s “defection” was not the first or arguably even the most important case of a U.S.-trained AI scientist relocating to China.

Harry Shum (沈向洋), Microsoft’s former executive vice-president and employee of the company for 23 years, emigrated to China in August 2020, one month before Zhu, to lead multiple Chinese AI research initiatives.⁷³ Pu Muming (蒲慕明), another committed AGI supporter whose Shanghai-area research empire rivals or exceeds Zhu’s, taught at the University of California, Irving; Yale; Columbia; the University of California, San Diego; and Berkeley for 36 years before giving up U.S. citizenship in 2017 and permanently returning to China.⁷⁴ Yao Chi-chih (姚期智), who taught AI and quantum computing at MIT, Stanford, and Berkeley and received the Turing Award in 2000, took up residence at Tsinghua University in 2004 and relinquished U.S. citizenship in 2016.⁷⁵

Nor can one ignore the fact that BIGAI is a Chinese undertaking through and through. This applies to physical infrastructure, intellectual capital, and, most distinctively, to its founding philosophy. Zhu argues that his project will use “Chinese philosophical thought” (中国哲学思想) as its foundational model, and goes to some length to link the Neo-Confucianism of Chinese philosopher Zhu Xi (朱熹, 1130–1200) and later interpretations of it to his vision of AGI.⁷⁶

On one level, this resembles the phrase “with Chinese characteristics” (以中国特色) that is appended to state-backed projects—which Zhu dutifully articulates—meant to instill ownership and pride.⁷⁷ But in this context, Zhu should be taken literally. Emphatically, *Chinese progress toward AGI should not be measured with an American or British yardstick.*

Zhu argues that China must stop “chasing the basketball all over the court” (追着篮球全场跑), meaning his country should stop following foreign paradigms, which he sees as a recipe for competitive failure, because technology fashions change. That is, China should “not blindly follow the current AI hotspots characterized by ‘big data, large computing power, and large models,’” but rather should “find a new way to explore our own path of scientific research and innovation with strong strategic determination.”⁷⁸

Zhu doubles down by arguing that China must think in terms of “you fight your way, and I fight my way” (你打你的, 我打我的), by which he means “move from a comprehensive passive defense to an active offense targeting key areas.”⁷⁹ He repeats the martial metaphor elsewhere, as in the “two bombs, one satellite” reference, in statements such as the following:

From the perspective of national security and economic development, this is an extremely important field, and it is also a must for China to participate in the international competition. It is a main battlefield.⁸⁰

Observations

General purpose AI (AGI) hypothetically confers an advantage on the organization or nation-state that first attains it. While the authors of this study do not see the contest in all-or-nothing terms, the concern is that the ramifications of this game-changing technology should fall under the control of political leaders whose ethical standards diverge widely from those of democratic systems.

China has been edging toward AGI since 2017—or, in the estimate of the authors of this study, earlier—when the concept was referenced in China’s “New Generation AI Development Plan.”⁸¹ Although AI developers Baidu Research, Alibaba’s DAMO Academy, and Tencent AI Lab on the commercial side; and the CAS Institute of Automation, Beijing Academy of AI, and sundry brain-inspired AI labs on the government side have research programs that may someday lead to AGI, the Beijing Institute is a watershed in China’s development of the technology, having named AGI as its sole, state-backed purpose.

Zhu’s pursuit of small data solutions to big tasks awaits a detailed description. Presumably, the plan is to replicate the brain’s ability to form and use abstractions, although it is unclear if this will be achieved through cognitive science (theories about the mind), neuroscience (the mind’s neural basis), or the two combined. BIGAI’s ability to coordinate its AGI program with other such efforts in China, including those that are

biologically oriented, or with projects linking in vivo neurons with computers, would warrant particular scrutiny.⁸²

The move puts China in principle on the same footing as the United States and UK, whose commercial enterprises are also working toward AGI, raising concerns about governance and safety. The United States and allied governments are encouraged to monitor these developments to support mitigation in the case of strategic challenges and collaboration in areas where all parties can benefit.

Authors

Huey-Meei Chang is CSET's senior China S&T specialist, co-editor of *Chinese Power and Artificial Intelligence: Perspectives and Challenges* (New York and London: Routledge, 2023), and co-author of several papers on China's technology acquisition and development. **William C. Hannas** is CSET's lead analyst and formerly the Central Intelligence Agency's senior expert for China open source analysis.

Acknowledgments

The authors of this study are grateful to John Chen, Igor Mikolic-Torreira, and Max Riesenhuber for their timely and useful critiques; to Annie Rehill for her thorough copyediting; Shelton Fitch for formatting the paper; and to Lynne Weil for overseeing preparation and publishing.



© 2023 by the Center for Security and Emerging Technology. This work is licensed under a Creative Commons Attribution-Non Commercial 4.0 International License.

To view a copy of this license, visit <https://creativecommons.org/licenses/by-nc/4.0/>

Document Identifier: doi: 10.51593/ 20230013

Endnotes

¹ “Exclusive | As a strategic scientist, Zhu Songchun, a top AI scholar, returned to China to prepare for the establishment of the Beijing General AI Research Institute” (独家 | 以战略科学家身份, 顶级 AI 学者朱松纯回国, 筹建北京通用 AI 研究院), Synced (机器之心), September 13, 2020,

<https://www.jiqizhixin.com/articles/2020-09-13>.

² James Barrat, *Our Final Invention: Artificial Intelligence and the End of the Human Era* (New York: Thomas Dunne Books, 2013).

³ See Seth Baum, “A Survey of Artificial General Intelligence Projects for Ethics, Risk, and Policy” (Global Catastrophic Risk Institute, Working Paper 17-1, 2017). Hugo de Garis’s AGI program at Xiamen University began in 2018.

⁴ William Hannas, Huey-Meei Chang, Daniel Chou, and Brian Fleeger, “China’s Advanced AI Research” (CSET, July 2022), 7–15, <https://cset.georgetown.edu/publication/chinas-advanced-ai-research/>.

⁵ “Toward General Artificial Intelligence” (迈向通用人工智能), 2, <https://www.bigai.ai/flip-book/>. BIGAI issued this 50-page color pamphlet in late 2022.

⁶ “Pursue the theory of grand unification, and be a colleague with Professor Zhu Songchun: Beijing Institute for General Artificial Intelligence sincerely invites talents from all over the world!” (追求大一统理论, 和朱松纯教授做同事: 北京通用人工智能研究院诚邀天下英才!), *AI Technology Review* (AI 科技评论), February 21, 2021, <https://posts.careerengine.us/p/6031e4cf607149081fe12f99>.

⁷ Beijing Institute for General Artificial Intelligence. See <https://www.bigai.ai/about/>.

⁸ The Wangxuan Institute of Computer Technology, formerly PKU’s Institute of Computer Science & Technology, was established in 1983 and subsequently renamed after its founder. Wangxuan Institute of Computer Technology: Introduction, https://www.icst.pku.edu.cn/english/about_us/introduction/index.htm.

⁹ “Toward General Artificial Intelligence,” 44.

¹⁰ This is a college within PKU where top students are allowed unprecedented freedom in choosing curricula and pace of advancement. It is named after former PKU president Cai Yuanpei (蔡元培).

¹¹ “Toward General Artificial Intelligence,” 44.

¹² “Toward General Artificial Intelligence,” 19. See also author attribution number 1 in Lifeng Fan, Manjie Xu, Zhihao Cao, Yixin Zhu, et al., “Artificial Social Intelligence: A Comparative and Holistic View,” *CAAI Artificial Intelligence Research* 1, no. 2 (2022): 144–60, which reads “National Key Laboratory of General Artificial Intelligence, Beijing Institute for General Artificial Intelligence (BIGAI),” <https://www.sciopen.com/article/10.26599/AIR.2022.9150010>.

¹³ Katja Grace, John Salvatier, Allan Dafoe, Baobao Zhang, et al., “Viewpoint: When Will AI Exceed Human Performance? Evidence from AI Experts,” *Journal of Artificial Intelligence Research* 62 (2018): 729–54.

¹⁴ Discussion with DeepMind’s top AGI cadre, September 2022.

¹⁵ Chinese AI research dates from the late 1970s but was not formally recognized by PRC funding agencies as a standalone discipline until 2018. William C. Hannas and Huey-Meei Chang, eds., *Chinese Power and Artificial Intelligence* (New York and London: Routledge, 2023).

¹⁶ Lex Fridman, “Origin of the term AGI (Ben Goertzel),” AI Podcast Clips, June 23, 2020, <https://www.youtube.com/watch?v=4X2xYyIk5x0>.

¹⁷ Grace, Salvatier, Dafoe, Zhang, et al., “Viewpoint: When Will AI Exceed Human Performance?”; William 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/>.

¹⁸ “Toward General Artificial Intelligence,” 1, 10.

¹⁹ “Toward General Artificial Intelligence,” 10, 11. Segments of the character are highlighted to spell out “A” “G” “I.”

²⁰ “Toward General Artificial Intelligence,” 21.

²¹ “Toward General Artificial Intelligence,” 19.

²² “Toward General Artificial Intelligence,” 25.

²³ “Toward General Artificial Intelligence,” 26.

²⁴ “Toward General Artificial Intelligence,” 27.

²⁵ “General Artificial Intelligence Experimental Class of Peking University (PKU ‘General Class’) 2021 Admissions Guide” (北京大学通用人工智能实验班 (北大“通班”) 2021 年招生简章), September 16, 2021, <https://yuanpei.pku.edu.cn/jyxx/jxxx/516434.htm>. PKU alone claims a staff of more than 100 AI instructors (教师).

²⁶ PKU CoRe Lab link, <https://pku.ai/>.

²⁷ “Toward General Artificial Intelligence,” 24-27.

²⁸ “Toward General Artificial Intelligence,” 1, 2, 9, 22.

²⁹ “2023 global campus recruitment starts! We invite you to go where no one has gone before and lead the new future of AI,” (2023 全球校园招聘开启！邀您一起探索无人之境，引领 AI 新未来), October 25, 2022, <https://www.bigai.ai/news/2023> 全球校园招聘开启！邀您一起探索无人之境，引/.

³⁰ “Pursue the theory of grand unification, and be a colleague with Professor Zhu Songchun: Beijing Institute for General Artificial Intelligence sincerely invites talents from all over the world!” The notice does not contain references to state or party loyalty, which is common in security-related job postings.

³¹ “Pursue the theory of grand unification, and be a colleague with Professor Zhu Songchun: Beijing Institute for General Artificial Intelligence sincerely invites talents from all over the world!”

³² Fan Lifeng's CV, <https://lifengfan.github.io/>; “Toward General Artificial Intelligence,” 45.

³³ Huang Siyuan's CV, <https://siyuanhuang.com>.

³⁴ Liang Yitao's CV, <https://www.ai.pku.edu.cn/info/1284/1977.htm>; “Toward General Artificial Intelligence,” 45.

³⁵ Liu Hangxin's CV, <https://liuhx111.github.io/>.

³⁶ Liu Tengyu's CV, <https://tengyu.ai/>.

³⁷ Peng Yujia's CV, “Toward General Artificial Intelligence,” 44.

³⁸ Yang Yaodong's CV, <https://www.yangyaodong.com/>.

³⁹ Zhang Muhan's CV, <https://muhanzhang.github.io/>.

⁴⁰ Zheng Zilong's CV, <https://zilongzheng.github.io/>.

⁴¹ Zhu Yixin's CV, <https://yzhu.io>.

⁴² “Focus interview: On the new journey to rejuvenate the country through science and technology the support of talent is still needed” (焦点访谈：新征程上科教兴国还需人才支撑), CCTV.com (央视网), November 25, 2022, <https://news.cctv.com/2022/11/25/ARTIWJN0iErt6666t22E1ACF221125.shtml>.

⁴³ “Focus interview: On the new journey to rejuvenate the country through science and technology the support of talent is still needed.”

⁴⁴ Spring courses 2022, <https://yzhu.io/courses/summer/22spring/>.

⁴⁵ Hannas, Chang, Chou, and Fleeger, “China’s Advanced AI Research,” 47–52.

⁴⁶ PKU and Tsinghua “AGI Class,” <https://www.bigai.ai/news/通班科研实训活动圆满结束/>.

⁴⁷ “Toward General Artificial Intelligence,” 3.

⁴⁸ Introduction to Zhu Songchun, <https://www.au.tsinghua.edu.cn/info/1110/2445.htm>.

⁴⁹ Introduction to Zhu Songchun.

⁵⁰ Introduction to “Distinctive Voices--Stuart Russell & Song-Chun Zhu ‘On Artificial General Intelligence and AI Alignment’,” February 13, 2023, <https://safeai.webs.upv.es/index.php/distinctive-voices-stuart-russell-songchun-zhu/>. Zhu had some 335 English language publications as of 2022, <https://vcla.stat.ucla.edu/publications.html>.

⁵¹ Zhu was selected for the program in 2010. “Zhu, Songchun” (朱松纯), Thousand Talents Project Net (千人计划网), June 9, 2014, <https://web.archive.org/web/20200524140234/http://news.jxnews.com.cn/system/2014/06/09/013148238.shtml>.

⁵² BAAI itself was established in 2018 under Huang Tiejun (黄铁军), vice-dean of PKU's Institute for Artificial Intelligence. Huang still holds this post within Zhu's BIGAI consortium. BAAI's goal is to integrate neuroscience, cognitive science, and information science as it builds “strong artificial intelligence” and “super-brain” (超脑) intelligent systems—goals wholly compatible with those held by Zhu. See Hannas, Chang, Chou, and Fleeger, “China’s Advanced AI Research.”

⁵³ The organization’s full name is Lotus Hill Institute for Computer Vision and Information Science (中国莲花山计算机视觉与信息科学研究院). “Songchun Zhu: Toward General Artificial Intelligence—From Big Data to Big Tasks” (朱松纯：走向通用人工智能——从大数据到大任务), Tencent Cloud (腾讯云), November 12, 2019, <https://cloud.tencent.com/developer/article/1537183>. Zhu is seated in the bottom row, third from right.

⁵⁴ In Zhu’s words, “In fact, we were the first team to do big data. In 2005, we led a group of international scholars, including those who later labeled ImageNet at Stanford and the later director of MIT’s lab.” See “Songchun Zhu: Toward General Artificial Intelligence—From Big Data to Big Tasks.”

⁵⁵ See “Songchun Zhu: Toward General Artificial Intelligence—From Big Data to Big Tasks.”

⁵⁶ “ChatGPT explodes, where is China's artificial intelligence going?” (ChatGPT 爆火，中国人工智能向何处去), ScienceNet.cn (科学网), March 6, 2023, <https://news.sciencenet.cn/htmlnews/2023/3/495386.shtm>. See also “Toward General Artificial Intelligence,” 2.

⁵⁷ “Stuart Russell's and Zhu Songchun's views on AGI and ChatGPT” (关于 AGI 与 ChatGPT，Stuart Russell 与朱松纯这么看), Synced (机器之心) February 28, 2023, <https://www.jiqizhixin.com/articles/2023-02-28-4>.

⁵⁸ “Stuart Russell's and Zhu Songchun's views on AGI and ChatGPT.”

⁵⁹ Irene Zhang, "AI Proposals at 'Two Sessions': AGI as 'Two Bombs, One Satellite'?", ChinaTalk, March 8, 2023, <https://www.chinatalk.media/p/ai-proposals-at-two-sessions-agi>. 两弹 is commonly mistranslated as "two bombs," i.e., atomic and thermonuclear weapons. The second 弹 (dàn) is actually an abbreviation of 导弹 (dǎodàn), "ballistic missile." See "Two bombs and one satellite' established China's status as a great power" ("两弹一星"奠定中国大国地位), CCTV.com (央视网), October 18, 2007, <http://discovery.cctv.com/special/C19607/20071018/104899.shtml>.

⁶⁰ Zhang, "AI Proposals at 'Two Sessions': AGI as 'Two Bombs, One Satellite'?"

⁶¹ "Zhu Songchun, member of the National Committee of the Chinese People's Political Consultative Conference: General artificial intelligence should be raised to the height of 'two bombs and one satellite'," (全国政协委员朱松纯: 应将通用人工智能提到“两弹一星”高度), *Southern Metropolis Daily* (南方都市报), March 6, 2023, <https://www.bigai.ai/news/nanfangdushibao/>.

⁶² Zhengwei Li (李政葳), "Zhu Songchun, Dean of PKU's School of Intelligence: Artificial intelligence will transform from weak artificial intelligence to general artificial intelligence" (北大智能学院院长朱松纯: 人工智能将从弱人工智能向通用人工智能转化), *Guangming Net* (光明网), September 22, 2022, https://edu.gmw.cn/2022-09/22/content_36042439.htm.

⁶³ According to Huang Tiejun: "Artificial superintelligence is the future—an evolutionary trend. Technological innovation and development should not be restricted by the limitations of human beings." At China National Computer Congress 2019 "Viewpoint forum 'Where is the moral boundary of artificial intelligence development?' was successfully held" (观点论坛「人工智能开发的道德边界在哪里?」顺利召开), October 24, 2019, <https://www.ccf.org.cn/YOCSEF/Branches/Shenzhen/News/2019-10-24/669602.shtml>.

⁶⁴ See Bio-sketch of Song-Chun Zhu, <http://www.stat.ucla.edu/~sczhu/bio.html>. One source states that the UCLA lab received "more than \$40 million in funding," while Zhu served as its principal investigator. "Decided! Professor Zhu Songchun, the master of AI vision, returned to China and became the dean of the Institute of Artificial Intelligence of Peking University" (定了! AI 视觉泰斗朱松纯教授归国就任北大人工智能研究院院长), *Sina Technology* (新浪科技), September 25, 2020, <http://tech.sina.com.cn/cs/2020-09-25/doc-iivhuipp6289115.shtml>.

⁶⁵ DARPA researcher's comments to the authors, June 2021: "He had a dramatic exit from the program last year . . . and we had to consider the security implications."

⁶⁶ "Focus interview: On the new journey to rejuvenate the country through science and technology the support of talent is still needed."

⁶⁷ Four of the UCLA lab's five faculty and many of its 28 Ph.D. students were affiliated with Chinese universities. Graduate students work in causal learning, cognitive modeling, communicative learning, generalization, grounding, reasoning and intent. See <https://vcla.stat.ucla.edu/people.html>; <https://ieeexplore.ieee.org/author/37281407500>.

⁶⁸ Tao Yuan, Hangxin Liu, Lifeng Fan, Zilong Zheng, et al., “Joint Inference of States, Robot Knowledge, and Human (False-)Beliefs,” IEEE International Conference on Robotics and Automation (ICRA), Paris, France, 2020, 5972–78.

⁶⁹ Luyao Yuan, Xiaofeng Gao, Zilong Zheng, Mark Edmonds, et al., “In situ bidirectional human-robot value alignment,” *Science Robotics* 7, no. 68 (July 13, 2022).

⁷⁰ Yixin Zhu, Tao Gao, Lifeng Fan, Siyuan Huang, et al., “Dark, Beyond Deep: A Paradigm Shift to Cognitive AI with Humanlike Common Sense,” *Engineering* 6, no. 3 (2020): 310–45, <https://www.sciencedirect.com/science/article/pii/S2095809920300345>.

⁷¹ See Center for Vision, Cognition, Learning, and Autonomy, “Milestones,” <https://vcla.stat.ucla.edu/news.html>; see also Bio-sketch of Song-Chun Zhu.

⁷² Bo Liu, Xidong Feng, Jie Ren, Luo Mai, et al., “A Theoretical Understanding of Gradient Bias in Meta-Reinforcement Learning,” accepted by NeurIPS 2022.

⁷³ See Harry Shum's (沈向洋) Baidu entry at <https://baike.baidu.com/item/沈向洋/10708178>; “Dr. Harry Shum was hired as the Director of the Academic Committee of Peking University's Institute of Artificial Intelligence” (沈向洋博士受聘北京大学人工智能研究院学术委员会主任), PKU website, August 1, 2020, <http://www.ai.pku.edu.cn/info/1086/1435.htm>.

⁷⁴ Pu Muming's CV (蒲慕明经历), http://www.ion.ac.cn/yjz/pmm_/ry/; “The Superstar of Science – Pu Muming” (科学界的超级明星——蒲慕明), SOHU.com, (搜狐), September 7, 2018, https://www.sohu.com/a/252558392_183834.

⁷⁵ See Andrew Chi-Chih Yao's (姚期智) biography at <https://iiis.tsinghua.edu.cn/en/yao/>; see also “Dialogue with Yang Zhenning and Yao Chi-Chih: Why did I give up my foreign nationality?” (对话杨振宁、姚期智：我为什么放弃外国国籍?), Xinhua.net (新华网), February 22, 2107, http://www.xinhuanet.com/politics/2017-02/22/c_129488963.htm.

⁷⁶ “Zhu Songchun: Interpreting ‘Ode to the Red Cliff’ from the perspective of artificial intelligence and talking about the balance between ‘heart’ and ‘reason’” (朱松纯：从人工智能的角度解读《赤壁赋》兼谈“心”与“理”的平衡), Synced (机器之心), January 7, 2022. <https://www.jiqizhixin.com/articles/2022-01-07-7>.

⁷⁷ Zhengwei Li, “Zhu Songchun, Dean of PKU's School of Intelligence:”

⁷⁸ Songchun Zhu, “Promoting original and leading innovation with organized scientific research” (以有组织科研推进原创性、引领性创新), *Guang Ming Daily* (光明日报), February 18, 2023, https://epaper.gmw.cn/gmrb/html/2023-02/18/nw.D110000gmrb_20230218_1-10.htm.

⁷⁹ Songchun Zhu, “Promoting original and leading innovation.”

⁸⁰ “Introduce top students to the AI field” (把顶尖学生引入人工智能领域), CCTV.net (央视网), April 28, 2021, <https://5gai.cctv.com/2021/04/28/ARTIQLganLFqyOik86neSq74210428.shtml>.

⁸¹ Hannas, Chang, Chou, and Fleeger, “China’s Advanced AI Research.”

⁸² Hannas, Chang, Chou, and Fleeger, “China’s Advanced AI Research”; and Hannas and Chang et al., “China’s Non-therapeutic Brain-Computer Interface Research” (CSET, forthcoming).