CSET CENTER *for* SECURITY *and* EMERGING TECHNOLOGY

China's Advanced AI Research

MONITORING CHINA'S PATHS TO "GENERAL" ARTIFICIAL INTELLIGENCE

智能科技

AUTHORS
Wm. C. Hannas
Huey-Meei Chang
Daniel H. Chou
Brian Fleeger

JULY 2022

CENTER *for* SECURITY *and* EMERGING TECHNOLOGY

Established in January 2019, the Center for Security and Emerging Technology (CSET) at Georgetown's Walsh School of Foreign Service is a research organization focused on studying the security impacts of emerging technologies, supporting academic work in security and technology studies, and delivering nonpartisan analysis to the policy community. CSET aims to prepare a generation of policymakers, analysts, and diplomats to address the challenges and opportunities of emerging technologies. CSET focuses on the effects of progress in artificial intelligence, advanced computing, and biotechnology.

CSET.GEORGETOWN.EDU | CSET@GEORGETOWN.EDU

JULY 2022

China's Advanced AI Research

MONITORING CHINA'S PATHS TO "GENERAL" ARTIFICIAL INTELLIGENCE

CENTER *for* SECURITY *and* EMERGING TECHNOLOGY

AUTHORS
Wm. C. Hannas
Huey-Meei Chang
Daniel H. Chou
Brian Fleeger

PRINT AND ELECTRONIC DISTRIBUTION RIGHTS

Cover illustration: "Intelligence science and technology" in Chinese.

DOCUMENT IDENTIFIER

doi: 10.51593/20210064

Contents

Executive Summary

This paper seeks to determine on the basis of publicly available information ("open sources") who in China is taking what steps toward *general artificial intelligence*, as shown by overt expressions and other common measures. While typically conceived as "artificial general intelligence" or AGI, this paper rejects that ambiguous term, along with its usual association with human-level machine intelligence, in favor of an approach that recognizes diverse pathways to broadly capable AI that functions autonomously in novel circumstances.

Accordingly, the paper examines what paths to *general AI* are available in principle, as a prelude to describing work underway in China to realize that capability. Three broad areas of Chinese research are identified as potentially germane: machine learning, brain-inspired AI research, and brain-computer interfaces (BCI). Data on the persons, institutions, and research making up this ecosystem is given as a foundation for downstream studies,* and as a starting point for a China-focused indications and warning watchboard.

The paper recounts the methodology used to build a database and prototype watchboard that enable analysts to capture and potentially forecast China AI-related events. Data supporting this pilot is conditioned to accept accretions from follow-on research, done locally or with outside participants, on Chinese artificial intelligence, AI's political uses, and *other emerging Chinese technologies*. The project aims to become or inspire a general foreign technology monitoring platform.

* Hannas and Chang, et al., "China's Cognitive AI Research," CSET (in preparation).

Introduction

In 2017, China's State Council released a comprehensive "New Generation AI Development Plan" aimed at making China the world's leading AI power by 2030.[1] While aspirational, China's Communist Party, national government, universities, research labs, and technology companies over the subsequent half decade have demonstrated an unwavering commitment to promote AI not only in applications, where China has a solid track record,[2] but also—as per the plan—in cutting-edge research that historically had not been China's forte.[3]

The plan has a complementary goal of achieving a "first mover advantage" (先发优势). In AI's case, this implies an ever-widening gap between the front-runner—China—and less capable nations, owing to AI's (projected) ability to generate better versions of itself. If China succeeds in either of these aims, the impact on the United States' and other liberal democracies' ability to compete may be severely tested. The risk extends beyond AI to technology in general due to AI's role as a technology *enabler.*

Accordingly, this paper explores China's efforts to make headway in three main sub-fields named in the 2017 plan: traditional (computational) AI research, "connectomics" or brain-inspired AI studies, and brain-computer interfaces. This report identifies institutions and persons that make up this ecosystem and provides descriptions of work done at some 30 organizations to convey a sense of the AI plan's actual scope and direction. China views this research as moving toward *general artificial intelligence*—software that can function in novel circumstances with greater autonomy and effectiveness—which may, or may not, resemble human intelligence.

Given the near absence of data-driven research on this topic—in Chinese language sources especially—by observers in the United States and Europe, and its potential import for competing nations, this paper also introduces an open-source pilot project, with extant examples, to monitor China's progress from narrow to broadly capable AI, using this study's data as a foundation. It concludes with a proposal to track China's AI growth as a target of concern and as a starting point for a science and technology (S&T) monitoring watchboard, which the United States currently lacks.

A study of China's technology development inevitably generates questions about state versus private management, foreign sources of inspiration, and impact on Chinese military power.[4] The authors address these topics in other studies to which the reader is referred. The short answers are: state involvement in Chinese AI is ubiquitous;[5] virtually all Chinese technology programs benefit from foreign know-how;[6] and Chinese planners are aware of AI's warfighting role.[7]

Research on next-generation AI with broad applications—unlike today's AI which is mostly limited to niche areas—leads to speculation that the outcome will be so-called "artificial general intelligence" (AGI), or by some accounts "artificial superintelligence," (ASI)—i.e., AI that matches or exceeds the intelligent behavior of humans in many or most respects. This paper does not discount that possibility nor treat it as especially probable, but proposes that advances in AI may be unpredictable or unrecognizable, taking forms that diverge from what laypersons regard as "intelligence" based on human experiences. While not necessarily the fabled holy grail of human-level machine intelligence, the outcome of China's advanced AI research will be as portentous.

Finally, describing advanced AI involves dissecting terms that interest AI professionals vitally and bear on how developments in this field are perceived. These sections (1.A and B) can be skipped by non-specialists—but at the cost of foregoing useful context.

The authors thank Dr. Dewey Murdick for helping frame the study, Helen Toner for serving as its "red teamer," Dr. Igor Mikolic-Torriera for his patience and encouragement, Dr. Catherine Aiken for her detailed review, Lynne Weil for marketing and production oversight, and Tessa Baker for administrative support. Thanks are also due Michael Miklaucic and John Chen for their reviews and substantive critiques, Maura McCarthy for her editorial review, Alex Friedland for copyediting support, and Ali Bours for layout and design.

1

General AI: Definitions and Pathways

Defining "AI" is problematic. In some ways it is a name in search of content. This section explores Chinese vs. U.S. and European views on AI's essence for a focal point from which to project its future trajectories.

A. RECTIFYING NAMES: "AGI" OR "GENERAL AI"[8]

AI has permeated global society and become a household word. Its successes notwithstanding, AI lacks a shared definition that captures what its makers are trying to create.[9] To the extent that the focus is on narrow tasks, definitions may be irrelevant—the AI executes its code, albeit through processes not wholly transparent, and a particular result is obtained. However, as scientists seek more capable AIs that function in broad domains without specific guidance, a better understanding of the goal—intelligence—is needed.

Defining "intelligence" was problematic even before computers.[10] Humans do not exhibit one type of intelligence but demonstrate cognitive skills across a spectrum of tasks, such as visual-spatial processing, quantitative reasoning, language use, and working memory—not to mention social and emotional intelligence. The discovery that *measures* of different mental abilities tend to correlate across tasks[11] led people to conceptualize human cognition in terms of a single "general intelligence factor," a term that is valid in a restricted sense but misleading,[12] and that has led to abuses.[13]

Moreover, measuring intelligence is not the same as understanding it. While it was possible for centuries to contemplate intelligence within a single biological domain, the advent of computing required an expanded definition to cover the problem-solving behavior of our electronic proxies. Inevitably, the sense that humans—as a seat of intellect—"own" intelligence, coupled with a tradition of psychological research, led to an overly human-centric view of what "intelligence" is. This anthropomorphic outlook obscured the need for a foundational inquiry into the "intelligence" part of artificial intelligence, to the detriment of the discipline and the ability to assess alternative AI programs—the discovery of which is a large part of any global monitoring regime.[14]

This complaint is widespread among specialists, not to mention analysts charged with informing and executing national policy, who struggle for a common definition of the terms of discourse.[15] Three leading artificial general intelligence experts explain how their niche has been impacted by terminological ambiguity:[16]

> "The debate on the essence of intelligence has been going on for decades, but there is still little sign of consensus."[17]

> "The lack of a clear, universally accepted definition is not unique to 'AGI.' For instance, 'AI' also has many different meanings within the AI research community, with no clear consensus on the definition. 'Intelligence' is also a fairly vague concept."[18]

> "For a long time, many different parties and factions in AI, adherent to more than one ideology, have been trying to build AI *without* understanding intelligence. And their habits of thought have become ingrained in the field, and even transmitted to parts of the general public."[19]

If "intelligence" and "AI" lack neutral (inclusive) definitions, "artificial general intelligence" is particularly susceptible to association with its human progenitor. AGI typically is defined as "the hypothetical ability of an intelligent agent to understand or learn any intellectual task that a human can," or "the capacity of an engineered system to display the same rough sort of general intelligence as humans," or "the representation of generalized human cognitive abilities in software."[20] Proposed measures to test AGI such as the Turing Test, Coffee Test, and Robot College Student Test similarly set their standard by comparison with a human model.[21]

This by itself is reasonable if human cognition is used as an extant example of a hypothesized *general problem solving ability* sought by software developers. The difficulty with casting this generalized goal in human terms is that the concept can be, and is, muddied by other human-like characteristics such as affect, consciousness, and morality,[22] causing AI's developers to aim past what is needed, lose sight of alternatives or, potentially, ignore the research of countries like China less wedded to anthropomorphic paradigms.[23]

Equating AGI with human cognition distorts the former and impoverishes the latter.[24] But the two are tied together so closely in lay understanding and even in specialized literature that the term AGI, in the authors' view, has lost its operative value.[25] Complicating matters, some approaches to AGI insist the AI be "embodied" human-like (have feedback mechanisms and reflect on its own outputs) to realize a general problem solving ability.[26]

If definitions matter within a language, they matter between languages. "Artificial general intelligence" is a recognized translation of the Chinese 通用人工智能 (tōngyòng réngōng zhìnéng), also rendered as "*general* (通用) *artificial intelligence* (人工智能)." Although both English terms are used in China, "general artificial intelligence" is preferred, as evidenced by its use in institutional nomenclature—e.g., the 北京通用人工智能研究院 (Beijing Institute for General Artificial Intelligence, BIGAI)[27]—and, of course, by the Chinese word order itself.

Herein lies the solution to the conundrum. By putting the adjective "general" *before* "artificial intelligence," China, possibly without thinking about it, solved in one shot the two problems previously identified. It eliminates the problematic assertion of "general intelligence" in its original context and human intelligence *per se* as a necessary goal for artificial emulation. What remains is a generalized form of artificial intelligence, difficult to achieve in its own right, but with no *necessary* connection to the mixed attributes attached to "AGI" outside China.

Use of the term "通用人工智能" in the hundreds of Chinese documents examined supports equating it with a *general search for more broadly capable artificial intelligence*, which may be comparable (by some measure) to human-level intelligence, or less capable, or exceed it in some or all domains. While the adjectives may be transposed ("人工通用智能"), it is far less common—a neologism inspired probably by the western "AGI," which this paper's authors hope does not take root. Use of the term "AGI" in this study—in cases where it cannot be avoided—is in the sense described here, namely "general AI" *sans* the anthropomorphisms and hyperbole.

Changing the term of reference to "general AI" comports with trends in AI's development from software with narrow to wide applications. An example is OpenAI's large language model GPT-3,[28] used in customer service, therapy, disinformation, and creative literature. Accordingly, this paper uses "general AI" as a synonym for "advanced AI." This formulation acknowledges brain-inspired AI research and artificially enhanced human intelligence as *parts* of AI's future but not necessarily the whole of it.

B. FIVE THEORETICAL PATHS TO GENERAL AI[29]

Having stated what this study is *not* about, a working definition of "intelligence" is needed. Three definitions taken from the literature follow (with comments from this report's authors in parentheses):

> "Intelligence measures an agent's ability to achieve goals in a wide range of environments."[30] (Narrow AI is an oxymoron.)

> "Intelligence is the capacity of a system to adapt to its environment while operating with insufficient knowledge and resources."[31] (Intelligence is generative and transferable.)

> "This definition sees intelligence as *efficient* cross-domain optimization."[32] (A system that depends continuously on big data and overspecification is not intelligent.)

These definitions, mutually compatible, reflect an understanding of AI's trajectory: narrow to wide, and allow a study of China's AI development without bias (advanced AI *requires* this or that computational approach), or framing the matter unproductively (who gets to "AGI" first).

The core question can now be addressed: what paths is China taking to enhance the cognitive power of agents under its jurisdiction or control?[33] Goertzel and Pennachin, in their foundational study of AGI's conceptual underpinnings, distinguish four types of approaches to AGI, each with extant programmatic examples, with and without links to human cognition:[34]

1. approaches that attempt to model biological brains
2. approaches explicitly guided by the human mind and brain
3. approaches inspired by the human mind much more than the brain
4. approaches that depend little on known science about human intelligence

This taxonomy is complemented by another schema, derived from a review of Chinese brain-inspired AI (BI-AI) programs, which qualifies "inspiration" as follows:[35]

- Inspiration by design. Accurate mathematical descriptions of real brain processes are run on computers to reproduce the behavior.
- De facto inspiration. Cognition is simulated by algorithms. The "inspiration" owes less to explicit modeling and more to equating ML successes with macro-level brain processes.

- Inspiration by default. Since "intelligence" derives originally from biological brains, efforts to emulate intelligence on artificial platforms are said to be brain-inspired.

Taking Goertzel and Pennachin's schema into account, the literature suggests there are (at least) five principled ways to build "advanced" AI:

1. Attempt to understand intelligence with cues from human behavior and create machine algorithms that emulate it.[36] This has been the majority viewpoint, associated with traditional ML/deep learning.[37]

2. Reverse-engineer a human brain on the assumption that what emerges is intelligence.[38] This "neuromorphic" or brain-imitative approach derives function from structure and is the province of "brain-inspired AI" and "connectomics."[39]

3. Force the emergence of intelligence, in theory, by running algorithms fast enough to "recreate the same amount of cumulative optimization power that the relevant processes of natural selection instantiated throughout our evolutionary past."[40]

4. Expand the definition of intelligence.[41] As we argue above, there is no reason to view intelligence as uniquely human.[42] Any "de novo" AI substantially able to achieve wide goals would qualify.

5. Finally, use brain-computer interfaces to position both elements, human and machine, to achieve (or overachieve) human goals. Embedded nanoscale chips and high-throughput cognitive "offloading" (partial brain emulation) are hypothetical approaches.

These categories are prototypes culled from the AI literature to illustrate possibilities. In practice, they merge and their boundaries are disputed. Yudkowsky, for example, treats whole-brain emulation as outside the AI family,[43] while Hansen views it as the paramount approach to AGI.[44] Baum sees WBEs as "computational entities with general intelligence" and hence within the AI pale,[45] and so on.

These paths tee up the study of China's approaches to advanced AI, while underscoring the fact that no one knows for certain how advanced AI will be realized; where and when it may happen; whether it will be smooth or discontinuous; or even what "it" is, as there are no precedents, no metrics to measure progress, and no clear break between AI and what meets (someone's) threshold for "advanced" AI.[46]

Accordingly, this paper examines three areas that cover some known possibilities, namely, so-called "mainstream" (主流) computational approaches to advanced AI; brain-inspired (类脑) AI research, including whole-brain emulation ("connectomics"); and brain-computer interfaces (脑机接口) that potentially result in cognitive augmentation. This division of paths to advanced intelligence is supported by most Chinese AI scientists, is grounded in state enactments dating from 2016, and sets the stage for effective mapping and monitoring (Sections 3 and 4 below).[47] However, it is possible that this paper's authors have overlooked some transformative development happening in a provincial lab or hiding in plain sight.

2 China's Advanced AI—A Seed List

The following is a record of China's top institutions engaged in one or more of the three typological areas of research identified as precursors to "general AI," namely, mainstream "computational" approaches, brain-inspired approaches, and brain-computer interfaces. The accounting is not exhaustive—that task is laid out in Section 3. Criteria for organizations to be included are professional status, accomplishments, published research, and declarations of intent known to the authors from Chinese internet sources judged to be reliable. Descriptions of research conducted at each are provided as an introduction to China's multiple approaches.[48]

C. CHINA'S MAINSTREAM APPROACHES TO ADVANCED AI[49]

Compute-intensive "big data" approaches to advanced (broadly capable) AI constitute both in China and globally the main focus of resources and attention. Here follow 10 such examples:

1. *Baidu Research* (百度研究院), technical arm of the Chinese search engine giant, is part of the company's Baidu AI Technology Platform, further divided into a cognitive computing laboratory (认知计算实验室), deep learning and big data labs, and others for biocomputing, robotics, safety, and quantum computing.[50] It is directed by Wang Haifeng (王海峰), Baidu's CTO. Baidu Vice President Wu Tian (吴甜) also serves as deputy director of China's National Engineering Laboratory for Deep Learning Technology and Applications (深度学习技术及应用国家工程实

验室),[51] which, under Baidu's leadership, announced in 2017 its intent to "build a research platform for general AI,"[52] complementing the NEL's work on computer vision, biometric ID, and human-machine interfaces.[53] In July 2021, Baidu partnered with Pengcheng Lab (鹏城实验室, C.6 below) to launch "ERNIE 3.0 Titan,"[54] a pre-trained large language model with 260 billion parameters[55] that achieved state-of-the-art scores on 60 natural language processing (NLP) tasks, demonstrating its ability to "generalize across various downstream tasks with a small quantity of labeled data."[56]

2. *Alibaba's DAMO Academy* (阿里巴巴达摩院), a branch of the e-commerce company, was founded in 2017 with facilities in Beijing, Hangzhou, and five sites abroad.[57] Its director is Zhang Jianfeng (张建锋), president of Alibaba Cloud Intelligence (阿里云智能). The academy hosts 16 laboratories researching everying from machine intelligence (机器智能) to quantum.[58] In June 2021, DAMO announced a trillion-parameter AI model with "nascent cognitive and creative abilities, and a goal of becoming the world's leading AGI (通用性的人工智能) model."[59] Five months later, it released a 10 trillion-parameter Multi-Modality to Multi-Modality Multitask Mega-transformer (M6) artificial intelligence system, touted as the world's largest AI model.[60] Zhou Jingren (周靖人), head of DAMO's data analytics and intelligence lab, explained that the goal is to improve M6 to "a level close to human beings" on the way to building general AI.[61] A 2022 document "Top Ten Technology Trends of the DAMO Academy" (达摩院十大科技趋势) regards "The ultra-large-scale pre-trained model as a breakthrough exploration from weak AI to general AI."[62]

3. *Tencent AI Lab* (腾讯人工智能实验室) is the research arm of Tencent (腾讯), the Shenzhen-based online gaming company and developer of the messaging app WeChat.[63] Founded in 2016, the lab researches computer vision, speech recognition, NLP, and ML.[64] It was under Zhang Tong's (张潼) tutelage from 2017 to 2019;[65] speech recognition expert Yu Dong (俞栋) became vice director in May 2017.[66] In 2019, Tencent's AI agent *Juewu* (绝悟) won recognition for defeating top international gamers,[67] while papers describing it were presented at the AAAI and NeurIPS conferences.[68] Tencent's goal of working toward "general AI" has been iterated in multiple venues. In July 2021, Juewu's development team acknowledged the software is "a seed (种子) for general artificial intelligence."[69] Tencent Vice President Yao Xing (姚星) claimed the company is "working together to tackle the ultimate goal of AGI."[70] And a July 2021 news release from Tencent's AI lab stated that "AI game research will be a key step for Tencent to overcome the ultimate AI research problem—General Artificial Intelligence (AGI)."[71]

4. *Huawei* (华为), the Shenzhen-based telecommunications equipment company, is heavily invested in AI, as demonstrated by its Ascend (昇腾) AI chips, Atlas AI computing platform used in "Safe Cities" applications, and MindSpore (昇思) AI computing framework.[72] The company signaled its formal entry into general AI development in May 2021, when it signed an agreement with the CAS Institute of Automation (CASIA, below) to build a "general artificial intelligence" platform in Wuhan.[73] In fact, Huawei engineers contemplated hybrid approaches to AGI as early as 2017[74] and continue to explore its computational costs and work-arounds.[75] The company's latest AGI precursor is Pangu α (盘古α),[76] a 200-billion parameter pre-trained language model developed with Peking University and Pengcheng Laboratory.[77] An indication of where Huawei may be headed is seen in its five-year AI investment plan, premised on achieving "general AI" by 2030—the projected date when China achieves world AI dominance[78]—and self-boot-strapping "artificial superintelligence" (超级人工智能) thereafter.[79]

5. *JD Research Institute* (京东AI研究院) was established in 2017 as a division of the Beijing-based e-commerce company JD.com.[80] The institute has three labs for ML, computer vision, and NLP aimed at achieving "human-like cognitive abilities" in language and speech.[81] This last goal is approximated by ViDA-MAN, "a digital-human agent for multi-modal interaction, which offers realtime audio-visual responses to instant speech inquiries" and sub-second latency.[82] JD Research Institute's leaders are a distinguished lot: Director of AI Research Zhou Bowen (周伯文) was chief scientist of IBM's Watson Group; He Xiaodong (何晓冬), JD's VP of technology, is a 15-year Microsoft veteran; Institute VP Mei Tao (梅涛) is another Microsoft alumnus. JD sponsors a second research arm—the JD Explore Academy (京东探索研究院)—set up in 2020 for trusted AI (可信人工智能), super deep learning, and quantum ML, all targeted at "disruptive" (颠覆式) innovation at the basic theoretical level.[83] The deep learning model reportedly works from insufficient data and is capable of knowledge distillation and transfer learning.[84]

6. *Pengcheng Lab* (鹏城实验室) was stood up in March 2018 as a Provincial Research Lab by the Guangdong and Shenzhen governments. The state-supported institution lists 59 universities, corporations, and research institutes as "strategic collaborators" beyond the Baidu and Huawei links discussed above.[85] Its outreach opportunities were further expanded in January 2019 with the establishment of a Pengcheng Lab International AI Development Center (鹏城实验室人工智能国际发展中心).[86] Originally a network IT research facility, AI has assumed an increasingly large part of the lab's portfolio. Pengcheng Lab has focused on data-intensive large AI models, where its ultra-fast Cloud Brain II (云脑2), developed

with Huawei to support AI computing,[87] was used to create Pangu α.[88] Although its director, Gao Wen (高文),[89] believes today's "narrow" AI will develop into "general AI,"[90] he is less confident traditional AI can achieve that prospect and is moving toward a BI-AI paradigm.[91] Pengcheng Lab's partnership with NEL-BITA (see D.9 below) in June 2019 is evidence of that transition.[92]

7. *Horizon Robotics* (地平线机器人) was founded in 2015 in Beijing's Haidian District by Yu Kai (余凯), an ML expert and former director of Baidu's IDL deep learning lab.[93] Chen Liming (陈黎明) became its CEO in 2021. The company specializes in AI chips for smart vehicles and computer vision. In 2018, Horizon established a "General AI Lab" in California's Silicon Valley, the only major Chinese company to research "general AI" there explicitly in name and in fact.[94] The lab is headed by Xu Wei (徐伟), a highly regarded deep learning expert, who left IDL because of "present AI's extremely numerous shortcomings" to produce machines "with learning abilities like humans."[95] Xu's goal to "form a small and sophisticated team focused on general artificial intelligence research" is reflected in job postings that seek "research scientists with strong background in fields of artificial intelligence.[96] Our mission is to build artificial general intelligence (AGI). We will focus on developing novel algorithms and technologies to allow machines to learn new knowledge and skills as efficient as humans."[97]

8. *Beijing Institute for General Artificial Intelligence* (北京通用人工智能研究院, BIGAI), under Ministry of Science and Technology and Beijing municipal auspices, aims to create a "grand unified theory" (大一统理论) of intelligence, a "general intelligent agent" (通用智能体), and is assembling the capacity to pursue these goals seriously.[98] The institute is led by returned UCLA professor and renowned AI scientist Zhu Songchun (朱松纯), in concert with PKU's Institute for Artificial Intelligence and Tsinghua's own (future) AGI institute.[99] Plans are for a staff of one thousand researchers drawn from China and "all over the world."[100] The authors of this report predict it will lead to clones, first in Shanghai then the provinces.[101] Beyond its scale, BIGAI is important for being the first state AI facility to bear the "AGI" name,[102] formalizing China's acceptance of a paradigm shift that we have been at pains to document,[103] from big data-dependent "narrow AI" to broadly capable AI that transfers learned patterns to new and unforeseen problems.[104] Its leadership views AGI as "the focus of international AI competition" in the coming decade.[105]

9. *CAS Institute of Automation* (中科院自动化研究所), established in Beijing in 1956, straddles "traditional" and "brain-inspired" AI research, exemplifying the shift from compute-heavy AI to broader, more practical models. CASIA is host to several AI luminaries, including Director Xu Bo (徐波), Zeng Yi (曾毅), who runs its

Research Center for Brain-inspired Intelligence (类脑智能研究中心), Jiang Tianzi (蒋田仔) head of its Brainnetome Center (脑网络组研究中心), and Tan Tieniu (谭铁牛), the computer vision expert and deputy chief of the PRC's liaison office in Hong Kong. Its multi-modal pre-trained model Zidong Taichu (紫东太初) with 100 billion parameters is seen as "an important first step from perceptual intelligence to general intelligence."[106] In late 2021, CASIA took the lead in a Multimodal Artificial Intelligence Industry Alliance meant to focus academic, laboratory, and industrial efforts on the emerging paradigm.[107] Meanwhile, the institute researches brain anatomical and functional connectivity at the micro-, meso-, and macro-scales and the application of these findings to advanced AI models.[108]

10. *Bohai University* (渤海大学) is AGI advocate Liu Kai's (刘凯) institute of affiliation and a platform for Temple University AI maven Pei Wang (王培).[109] Professor Wang is a world-class AGI innovator, second in name recognition only to Ben Goertzel, whose multiple publications decisively shaped our thinking.[110] Wang's non-axiomatic model of intelligence offers a uniform foundation for AI and human thought and support for a credible AGI model—NARS.[111] Liu Kai chairs regular sessions of the China Artificial General Intelligence Annual Conference (中国通用人工智能年会), an extension of the global event that Goertzel and Wang tend to dominate; the China sessions are followed by AGI workshops.[112] The 2021 convention featured keynote talks on number sense, computer vision, and methods for building self models.[113] Liu, who taught at Wuhan's Huazhong Normal University (华中师范大学) until 2017, teaches machine education (机器教育), computational psychiatry, and brain-like systems in the university's College of Education Science, an uncommon affiliation that makes perfect sense in an AGI context.[114]

BOX 1

Major Chinese entities pursuing general AI via traditional (computational) approaches

- *Baidu Research*
- *Alibaba's DAMO Academy*
- *Tencent AI Lab*
- *Huawei*
- *JD Research Institute*
- *Pengcheng Lab*
- *Horizon Robotics*
- *Beijing Institute for General Artificial Intelligence*
- *CAS Institute of Automation*
- *Bohai University*

D. BRAIN-INSPIRED AI, CONNECTOMICS, WHOLE-BRAIN EMULATION

Chinese critics of mainstream "big data–small task" computational approaches to general AI point to biological brains as models of inspiration or, in extreme cases, detailed emulation. The extent of such Chinese efforts to "merge" artificial and biological intelligence can be discerned in the following examples:

1. The *Beijing Academy of Artificial Intelligence* (北京智源人工智能研究院), also called the Zhiyuan Research Institute (智源研究院), was established in 2018 under Huang Tiejun (黄铁军), vice-dean of PKU's Institute for Artificial Intelligence (人工智能研究院),[115] to integrate neuroscience, cognitive science, and information science, on its way to building "strong artificial intelligence" (强人工智能) and "super-brain" (超脑) intelligent systems.[116] China's Ministry of Science and Technology described the goal as "transformative and disruptive breakthroughs."[117] Xu Bo (徐波), Tang Jie (唐杰), and Liu Jia (刘嘉), all strong proponents of AGI, have leadership roles.[118] Zhang Hongjiang (张宏江), former CTO of Microsoft Research Asia is also on its board. Recent accomplishments are the Shenji (神机) series of simulation platforms and a "Bio-intelligence Open Source Platform" (生物智能开源开放平台) with five components.[119] BAAI is the home of Wudao (悟道) 2.0, a multi-modal AI model comparable to GPT-3, meant to "enable machines to think like humans and move toward general AI."[120]

2. Beijing Normal University's (BNU) *State Key Laboratory of Cognitive Neuroscience and Learning* (认知神经科学与学习国家重点实验室) was established in 2005 to perform micro- and mesoscopic connectomics research, with a recent focus on reward processing and long-term episodic memory.[121] BNU is one of three Chinese universities, with Peking and Tsinghua, allied with MIT's McGovern Institute for Brain Research.[122] The same three universities and a fourth Beijing facility—the CAS Institute of Psychology (中科院心理研究所)—form the Chinese end of a "Transregional Collaborative Research Centre on Crossmodal Learning" run by Zhang Jianwei (张建伟) and colleagues at Universität Hamburg, aimed at describing "the neural, cognitive, and computational mechanisms of crossmodal learning."[123] Its goals are improved deep learning and use of brain-computer interfaces to "accelerate AI."[124] The CAS Institute of Psychology is home to Zuo Xinian (左西年), lead author of "An open science resource for establishing reliability and reproducibility in functional connectomics."[125] Zuo's doctorate is from BNU.[126]

3. The *Chinese Institute for Brain Research* (北京脑科学与类脑研究中心, CIBR), also in Beijing, dates from March 2018. CIBR was "strategically deployed" by the

city's S&T commission as a cooperative framework for Beijing-area universities, the PLA Academy of Military Science (军事科学院), and others.[127] Its mission is "co-ordinating research institutes and managing research programs under the guidance of the China Brain Initiative and Beijing Brain Initiative, and making Beijing the world epicenter for neuroscience and brain-inspired computation."[128] Besides medical research, CIBR studies brain-inspired AI, optical imaging, and brain-computer interaction (脑机交互作用).[129] The institute is co-directed by Rao Yi (饶毅), president of Capital Medical University in Beijing, former dean of sciences at Peking University, and founding director of the PKU-IDG/McGovern Institute for Brain Research, and Luo Minmin (罗敏敏), an investigator at Beijing's National Institute of Biological Sciences and a professor at Tsinghua University. Pu Muming and the Allen Institute's Christof Koch are on its advisory board.[130]

4. The *State Key Laboratory of Brain & Cognitive Science* (脑与认知科学国家重点实验室), another Beijing-based institute, was established in 2005 as a part of CAS's Institute of Biophysics (生物物理研究所). The lab does multi-disciplinary research on the "cognitive basic unit," learning and decision-making, and neural mechanisms of information processing in drosophila and non-human primates,[131] supported by an "ultra-high field MRI" platform used for brain imaging.[132] It is headed by Chen Lin (陈霖), who specializes in visual cognition and brain imaging. Its academic committee reads like a who's who in China BI-AI, including BAAI's Xu Bo; CIBR's Luo Minmin; Guo Aike (郭爱克), also at CAS's Institute of Neuroscience (CAS-ION, functional brain mapping, mind-body problem); Zhang Xu (张旭), director of the Chinese Academy of Science and Technology's "Neuroscience Direction Forecasting and Technology Roadmap project;" and CASIA's Tan Tieniu (谭铁牛).[133] The lab is creating a platform to track visual cognition from the genetic to behavioral level.[134]

5. CAS's *Center for Excellence in Brain Science and Intelligence Technology* (脑科学与智能技术卓越创新中心, CEBSIT), established in 2014, is an umbrella organization for 39 research institutions distributed in Beijing, Shanghai, and 13 other locations.[135] The CAS Institute of Neuroscience (神经科学研究所, ION) and CASIA are its main "supporting units."[136] CEBSIT is directed by Pu Muming (蒲慕明, AKA Muming Poo), who also manages CAS-ION.[137] Its vice directors are Xu Bo and Du Jiulin (杜久林), a neuroscientist and ION's vice director. Tan Tieniu of CASIA is its chief scientist. CEBSIT researches whole brain connectomics (全脑联结组), multi-sensory mode perception, computational models for semantic comprehension, and neuron-inspired computing chips."[138] In 2018, the center established a "G60 Brain Intelligence Innovation Park" (G60脑智科创基地) under Shanghai

municipal auspices[139] with a USD $1.5 billion budget for BI-AI.[140] The facility uses cloned monkeys to eliminate variables between specimens.[141] Pu Muming has been running connectome projects here since the turn of the century.[142]

6. Fudan University's *Institute of Science and Technology for Brain-Inspired Intelligence* (复旦大学类脑智能科学与技术研究院, ISTBI), launched in 2015 in Shanghai, hosts "centers" for cognitive neuroscience, computational biology, big data biomedical science, biomedical imaging, neural and intelligent engineering, and brain-inspired chips.[143] Run by Feng Jianfeng (冯建峰),[144] it boasts the world's largest brain science database, with access to the U.S. Human Connectome Project, the UK's Biobank, ISTBI's own 10 terabyte holdings, and the "largest number of magnetic resonance imaging devices in Asia" paired with AI algorithms to screen the images.[145] ISTBI's Zhangjiang International Brain Imaging Center (张江国际脑影像中心) reportedly "is building the world's most advanced and Asia's largest ultra-high-end scientific research magnetic resonance system." The center claims to have built "internationally leading intelligent algorithms and spatio-temporal data analysis and processing software" and "the world's largest full-dimensional brain database and algorithm center."[146]

7. Shanghai Jiao Tong University's *Center for Brain-like Computing and Machine Intelligence* (上海交通大学仿脑计算与机器智能研究中心, BCMI) was founded in 2002. Its mission, described in 2013, is "to understand the mechanism of intelligent information processing and cognitive process in the brain" with research in computer vision, NLP, cognitive computing, BCIs, and electroencephalography (EEG) signal processing.[147] A 2021 skill search by the center listed scene perception and understanding, commonsense learning, multimodal interaction learning, selective attention, and visual causal reasoning—all indicative of AI's evolution at BCMI and in China.[148] The center's Lu Baoliang (吕宝粮, E.6 below) and Zhang Yaqian (张亚倩) have published on "affective BCI" that "can recognize and modulate human emotion" and its role in general AI—a technology that has raised ethical concerns in some circles.[149] The center is complemented by SJTU's Artificial Intelligence Research Institute, Machine Cognitive Computing Research Center (人工智能研究院, 机器认知计算研究中心).[150]

8. The *Shanghai Center for Brain Science and Brain-inspired Intelligence* (上海脑科学与类脑研究中心, BSBII)—or "Shanghai Brain/AI Center"—was set up by CAS in 2018,[151] the same year its Beijing-based namesake, CIBR (see above), was established. BSBII is a coordinating center for Yangtze Delta BI-AI research and is linked with CEBSIT, the Fudan and Shanghai Jiao Tong University institutes, the

HUST-Suzhou Institute for Brainsmatics, and other regional players.[152] Organizationally, it is part of the Zhangjiang Lab's Institute of Brain-Intelligence Technology (张江实验室脑与智能科技研究院, BIT), stood up a year earlier in Pudong's Zhangjiang Science City and run by Zhang Xu (张旭).[153] BSBII itself is part of Pu Muming's local empire; Feng Jianfeng and Du Jiulin are vice directors.[154] Both BIT and the BSBII center collaborate with iFlytek and other Chinese AI companies.[155] BSBII operates BCI and "brain atlas big data" platforms. Research includes macro- and mesoscopic connectomics, BI-AI, and brain-inspired computing devices.[156]

9. The *National Engineering Laboratory for Brain-inspired Intelligence Technology and Application* (类脑智能技术及应用国家工程实验室, NEL-BITA) was established in Hefei's High-tech Zone under the University of Science and Technology's (中国科学技术大学) auspices in May 2017.[157] A member of the AI Industry Technology Innovation Strategic Alliance,[158] its research priorities are brain cognition and neural computing, brain-inspired multimodal sensing, brain-inspired chips, quantum AI, and brain-inspired intelligent robots.[159] NEL-BITA's director is Wu Feng (吴枫), former chief researcher at Microsoft Research Asia and assistant dean of USTC. Zha Zhengjun (查正军) is its executive director; Sun Xiaoyan (孙晓艳), another MSRA alumna, is deputy director. Although Pu Muming is on its board, none of the other 23 board members or its top research staff are among those persons mentioned above. In September 2017, the lab spun off a commercial venture to integrate its AI technologies in computer vision, "small sample learning" (小样本学习), and cross-media multimodal analysis with industry.[160]

10. The *HUST-Suzhou Institute for Brainsmatics* (华中科技大学苏州脑空间信息研究院) was established in 2016 at Huazhong University of Science and Technology (in Wuhan). Its director is Li Pengcheng (李鹏程) of HUST's National Laboratory for Optoelectronics (武汉光电国家实验室).[161] Deputy directors are Gong Hui (龚辉) and Li Anan (李安安). Notable researchers are Ye Chaohui (叶朝辉) and Luo Qingming (骆清铭), both CAS academicians. Ye was director of CAS's Wuhan branch and has held high posts at Wuhan institutes.[162] Luo founded Brainsmatics and is president of Hainan University.[163] The institute uses micro-optical sectioning tomography to model a high-resolution mammalian brain.[164] Based on "structural and functional imaging of neuron types, neural circuits and networks, neural-glial interfaces, vascular networks, etc. with high temporal-spatial resolution and specific spatial locations," Brainsmatics is working to decipher brain function and "promote brain-inspired artificial intelligence by extracting cross-level and multi-scale temporal-spatial characteristics" of brain connectivity.[165]

BOX 2

Major Chinese entities pursuing general AI via brain-inspired research

- *Beijing Academy of Artificial Intelligence*
- BNU's *State Key Laboratory of Cognitive Neuroscience and Learning*
- *Chinese Institute for Brain Research*
- *State Key Laboratory of Brain & Cognitive Science*
- CAS's *Center for Excellence in Brain Science and Intelligence Technology*
- *Fudan University's Institute of Science and Technology for Brain-Inspired Intelligence*
- *Shanghai Jiao Tong University's Center for Brain-like Computing and Machine Intelligence*
- *Shanghai Center for Brain Science and Brain-inspired Intelligence*
- *National Engineering Laboratory for Brain-inspired Intelligence Technology and Application*
- *HUST-Suzhou Institute for Brainsmatics*

E. BRAIN-COMPUTER INTERFACES AND NEUROMORPHIC CHIPS[66]

Brain-computer interfaces use AI to improve their operation, while opening a path to cognitive enhancement. Neuromorphic chips that imitate brain structure promise faster processing speeds for algorithms that support general intelligence and are being adapted in China for use in BCIs. Examples of both are provided.

1. *Tsinghua University BCI Lab* (清华大学脑机接口研究组) was established in 2004 and is one of several interlocking labs and institutes operated by the university, including also Tsinghua University Institute for Brain and Cognitive Sciences (清华大学脑与认知科学研究院, THUIBCS), run by Dai Qionghai (戴琼海),[167] and Tsinghua Laboratory of Brain and Intelligence (清华大学脑与智能实验室), whose research includes computational neuroscience and BI-AI.[168] Liu Jia (刘嘉, D.1 above), Zhu Jun (朱军), and Gao Xiaorong (高小榕) are prominent members.[169] Gao runs the BCI Lab with colleague Gao Shangkai (高上凯) and is its leading figure.[170] The lab studies the application of BCI to cognitive skill assessment, which besides med-

icine can also be used in lie detection and "human-machine collaboration."[171] It is also reportedly pushing the frontier on high-throughput interfaces, including wireless BCI transmission.[172] The lab was the first to implement non-invasive BCI technology based on steady-state visual-evoked potential (SSVEP), used to study the relationship between physical stimuli and human cognition.[173]

2. *The Center for Brain-Inspired Computing Research* (脑计算研究中心) was established by Tsinghua University in 2014 to study neural functional/computational theory, machine learning, and chip architecture.[174] The center draws on faculty from seven departments in "brain science, electronics engineering, microelectronics, computer science, automation, materials science, and precision instruments."[175] In 2019, a team led by Center Director Shi Luping (施路平) announced a brain-inspired computing chip called "Tianjic" (天机 Chip) and software tool chain able to "simultaneously support the neural network models of computer science and neuroscience,"[176] thus providing a platform for AGI.[177] The chip is partly analog and partly digital and meant to mimic the computational principles of biological brains. A year later they built a Turing-complete software model to bridge the divide between traditional "computer-science-based artificial neural networks" and neuroscience-driven AI models.[178] Co-developer Pei Jing (裴京) and Shi Luping market the chip through Beijing Lingxi Technology Co. (灵汐科技).[179]

3. *NeuraMatrix* (宁矩科技) was incubated by Tsinghua University in 2019 to build "active implantable systems interfacing with the human body and artificial devices."[180] As its name suggests, the project draws inspiration from the U.S. company Neuralink.[181] China has decades of experience developing non-invasive BCI systems, but this is the country's first effort to create an implantable device.[182] Moreover, unlike other Chinese BCI projects whose spoken goal is to alleviate disability, NeuraMatrix states openly its aim to augment the cognitive power of healthy persons "by effectively merging human and artificial intelligence,"[183] fulfilling the 2017 "New Generation AI Development Plan" to effect a merger, misinterpreted by some as metaphor.[184] The full-service package will include electrode materials, a neural interface chip, "infinite multi-point interface equipment, a signal acquisition and analysis platform and system-level brain-computer interface platform."[185] NeuraMatrix's founders are Bai Shuo (白硕, CAS) and Zhang Milin (张沕琳, Tsinghua's Laboratory of Brain and Intelligence, E.1 above).[186]

4. Tianjin University's *Brain Science and Brain-like Research Center* (科学与类脑研究中心) was established in September 2019.[187] It is complemented by the university's Institute of Medical Engineering and Translational Medicine (医学工程

与转化医学研究院) and a Neural Engineering Center (天津神经工程中心) focused on brain cognition, medicine, and BCI.[188] All three units are directed by Professor Ming Dong (明东).[189] In July 2019, prior to the center's establishment, the university unveiled its "Brain Talker" (脑语者) chip able to separate signal from noise with great accuracy.[190] The chip will "replace traditional computer devices used in BCI" thanks to its portability, greater precision, and faster transmission rates.[191] This non-invasive system was developed jointly with the state-owned China Electronics Corp. (中国电子信息产业集团).[192] Xu Minpeng (许敏鹏), assistant director of the Brain Science Center and project lead, acknowledges China's gap in invasive BCI but claims its non-invasive technology is "world class." A follow up chip reportedly three times faster than competing systems is under development.[193]

5. Fudan University's *Institute of Brain-inspired Circuits and Systems* (类脑芯片与片上智能系统研究院) was set up in July 2017 to support the China Brain Project and Shanghai's Zhangjiang National Lab (张江国家实验室) with research on BI-AI chips and neuron signal acquisition.[194] It is led by Min Hao (闵昊), an expert on non-volatile memory and wireless chip design.[195] In 2021, the institute announced China's first wireless BCI circuit for transmitting information between a chip and nerve cells. Team leader Ye Dawei (叶大蔚) claims it "outperforms foreign versions on many levels" at half the cost.[196] The Ministry of Education's Key Laboratory of Brain Functional Genomics (脑功能基因组学教育部重点实验室) at Huadong Normal University supported its development.[197] The device is currently installed on the skulls of freely moving mice. Its technical aspects are described in a paper by eight of the institute's researchers,[198] who also claim affiliations with the State Key Lab of ASIC and Systems (专用集成电路与系统国家重点实验室), which has been operating at Fudan University since 1995.[199]

6. Shanghai Jiao Tong University's *Ruijin Hospital BCI and Neuromodulation Center* (瑞金医院脑机接口及神经调控中心) was founded in 2020 with a goal of using BCI to address depression and other types of neuropsychological illnesses—BCI's "affective" applications.[200] The center is co-directed by Sun Bomin (孙伯民) and Lu Baoliang, the latter co-posted to the Center for Brain-like Computing and Machine Intelligence (D.7 above). According to Professor Sun, the plan is to "implant chips into patients' brains via a minimally invasive surgery" followed by electrical stimulation based on AI analysis.[201] Meanwhile, Lu in December 2021 spun off a BCI enterprise in partnership with Chinese video game giant miHoYo (米哈游) called Língwéiyīsī (零唯一思, no English name),[202] which—curiously—focuses on medical research (including BCI) *and* game development.[203] MiHoYo, for its part, in March that year agreed to build a laboratory with Ruijin Hospital's

Encephalopathy Center (脑病中心) run jointly by Lu and miHoYo's Anti-Entropy Studio (逆熵工作室) to develop brain interface technology.[204]

7. Zhejiang University's *Frontier Science Center for Brain and Brain-Machine Integration* (脑与脑机融合前沿科学中心), also called the Double Brain Center (双脑中心), was established under the Ministry of Education's auspices in October 2018. Wu Zhaohui (吴朝晖) is director and Duan Shumin (段树民) is its chief scientist.[205] Duan doubles as dean of Zhejiang's University's School of Brain Science and Brain Medicine (脑科学与脑医学学院), built in 2019 to research hybrid intelligence, BCI, and brain-inspired computing.[206] There is also a Zhejiang Laboratory (之江实验室), stood up in 2017 under provincial Communist Party auspices to guide development of a "national strategic scientific and technology force" in AI areas.[207] Zhejiang University's research in BCI dates from 2012. It has also achieved a number of BI-AI "firsts," including the world's largest neuromorphic computer in 2020 called the "Darwin Mouse."[208] The device, built by Zhejiang Lab, contains 792 "Darwin II" BI chips that can emulate around 120 million spiking neurons and 100 billion synapses—about the same as a mouse—on just 350–500 watts.[209]

8. South China University of Technology's *Center for Brain Computer Interfaces and Brain Information Processing* (脑机接口与脑信息处理中心) was established in 2007 in Guangzhou to research brain-computer interaction and large-scale brain data.[210] It is directed by Li Yuanqing (李远清),[211] who also runs the university's School of Automation Science and Engineering (自动化科学与工程学院). The center has five research teams whose topics collectively cover EEG and fMRI signal analysis, sparse signal representation, pattern recognition, ML, neural networks, big data processing, robotics, and, of course, BCI. Li acknowledges that in China BCI research "is still mainly aimed at normal people (正常人为主)."[212] In March 2019, the center formed a joint venture with China AI giant iFlytek (科大讯飞) called South China Brain-computer Interface Technology (华南脑控智能科技, iHNNC),[213] founded by Li and managed by CEO Xiao Jing (肖景).[214] The company sees BCI as empowering "the fields of elderly care and disability, mental and spiritual health, education, entertainment, security, military and other fields."[215]

9. *Tianqiao and Chrissy Chen Institute* (陈天桥雒芊芊研究院, TCCI) was founded at Caltech in 2016 by online gaming pioneer Chen Tianqiao (陈天桥) and spouse Chen Qianqian (陈芊芊). In 2020 and 2021, the institute established two "Frontier Labs" (前沿实验室) for brain research, one at Huashan Hospital (华山医院) in Shanghai, the other at the Shanghai Mental Health Center (上海精神卫生中心) whose vision statement reads: "To enrich people's lives with brain- and mind-related

technology."[216] The institute intends to extend development of its invasive BCI products beyond therapeutics to cognitive augmentation[217] and hopes eventually to address scientifically every important mind-body problem in the history of philosophy, from the cognitive basis for beliefs through free will, emotion, AGI, and cognitive uploads aimed at immortality.[218] Tao Hu (陶虎), who runs the Huashan facility, in late 2021 founded his own BCI company NeuroXess (脑虎科技) to build invasive BCI that can "twin" (孪生) human and artificial intelligence.[219] Tao is joined by billionaire Peng Lei (彭蕾, Lucy Peng), one of Alibaba's co-founders.[220]

10. *CAS Institute of Automation* (中科院自动化研究所, CASIA) besides traditional and BI-AI research is also invested in BCI development, and its importance to China's Brain-AI program warrants this double listing. CASIA's work on coding and decoding of visual neural information, see as "the core technology of brain-computer interfaces,"[221] is regarded as "an important stepping stone in the work to create better brain-machine interfaces."[222] CASIA researcher Yu Shan (余山), who is also deputy director of China's State Key Laboratory of Pattern Recognition (模式识别国家重点实验室副主任) with specialties in brain information processing, BI-AI, and BCI,[223] sees BCI "ultimately enhancing and expanding brain functions," which puts CASIA into the augmentation camp.[224] A 2021 report titled "BAAI AI Frontiers" listed among CASIA's achievements a robotic system that can "accurately implant flexible electrodes into the cerebral cortex of animals under the guidance of microscopic images," thus laying a path toward invasive BCIs.[225]

BOX 3

Major Chinese entities pursuing general AI via brain-computer interfaces or neuromorphic chips

- *Tsinghua University BCI Lab*
- Tsinghua University's *Center for Brain-Inspired Computing Research*
- *NeuraMatrix*
- Tianjin University's *Brain Science and Brain-like Research Center*
- Fudan University's *Institute of Brain-inspired Circuits and Systems*
- Shanghai Jiao Tong University's *Ruijin Hospital BCI and Neuromodulation Center*
- Zhejiang University's *Frontier Science Center for Brain and Brain-Machine Integration*
- South China University of Technology's *Center for Brain Computer Interfaces and Brain Information Processing*
- *Tianqiao and Chrissy Chen Institute*
- *CAS Institute of Automation*

3 Monitoring for Safety and Security

This section makes a case for monitoring China's AI development as a bellwether for AI risk, understood in terms of global *safety* and U.S. national *security* (both 安全 in Chinese).[226] The authors' efforts to build a relational database to support discovery and implement a scalable indications and warnings (I&W) watchboard are described.

F. WHY BUILD A CHINA AI "WATCHBOARD?"

In terms of global AI safety, China is one of many countries—the United States included—able to create advanced forms of artificial intelligence that pose unknown and potentially catastrophic risks. From a national security standpoint, China is no more or less likely than a half dozen other states to inflict harm on the United States and its allies through advanced AI when motivation is considered. So why focus on China?

There are multiple reasons why U.S. policymakers may want to support a rigorous monitoring regime able to track China's advances toward general AI:

China intends to lead the world in AI and achieve a "first mover" advantage.

China's "New Generation AI Development Plan," released in 2017, and other official proclamations declare China's intent to lead the world in artificial intelligence by 2030 and achieve a first mover advantage (先发优势) through general AI.[227] These goals, if realized, have serious implications for U.S. security, since early success at building advanced AI portends not only commercial and strategic advantages but potentially an asymptotic accretion of cognitive resources (a so-called "intelligence explosion") widely viewed as an existential risk and game-changer vis-à-vis competing nations.[228]

Data-based monitoring will guard against hype and overreaction to imagined threats.

Fair-minded patriots note these proclamations, take China at its word, imbibe fear-mongering hype from pundits misinformed by the same sources—and support a crash program that leads to a dangerous and unnecessary "AI arms race." This scenario is not without precedent (the 1950-60s alleged U.S. "missile gap" with the Soviet Union) and, in the AI realm, has nearly happened at least once already.[229] Accurate and timely accounting of China's progress toward advanced AI will reduce the likelihood of miscalculation.

China's ownership and export of general AI will exacerbate existing political abuses.

Risks from weaponized or flawed AI is not the only cause for concern. Equally worrisome is China's use of AI for political oppression and its ability to export this technology through its market power. The problem is documented by CSET's Dahlia Peterson[230] and by Maya Wang at Human Rights Watch.[231] Also, credible cognitive neuroscientists raise the specter of AI-based technologies facilitating mind control in service of totalitarian goals—not as a metaphor ("influence operations") but literally at the neuron level.[232] Preventing this implies an ability to track its development.

China has a poor track record in safety and post-disaster communication.

It is hard to ignore a near-existential calamity that originated in China. Wherever one puts the locus of the COVID-19 pandemic, two datapoints emerge: a massive scourge began in China, despite safeguards, and China responded lethargically to global calls for information. Can one assume an AI disaster will be treated differently? There is little chance China will allow direct safety inspections of its AI infrastructure or eschew technology that leads to national advantage, hence the need for a comprehensive watchboard able to track unsafe developments.

Data-driven knowledge leads to opportunities for collaboration on safety.

Although one does not usually think in these terms, there is a need for accurate and timely data to alert U.S. and allied policymakers for opportunities to engage China on AI matters, not only as they involve safety, but also in areas where both countries and the world stand to benefit, such as AI algorithms that can pinpoint inputs to climate change, promote species preservation, and so on. Spotting these opportunities requires timely monitoring and communication, which is not fully possible under present circumstances.

China is already systematically monitoring U.S. AI developments.

Watchboards can be viewed as unfriendly acts, akin to espionage when conducted by governments, and serve as grist for United Front propaganda that portrays China as a "victim" of racism or foreign aggression. In fact, China has operated an open-source science and technology intelligence (科技情报, STI) network since 1958 that is mainly U.S. focused, outstrips U.S. efforts by two or three orders of magnitude, and is immeasurably more effective.[233] The present project does not redress the imbalance but may lead to the U.S. government accepting a serious and badly needed open-source-monitoring solution.

Monitoring will alert U.S. authorities to technology transfers not in America's interest.

One argument for expanded open-source monitoring, proposed in various forms and forums,[234] is that it is impossible to dissociate China's indigenous research (its own business) from predatory technology transfers (the world's business). The same applies here—any assessment of China's AI progress must include an accounting of what it obtains from abroad. Transactions captured by the watchboard inimical to U.S. national and economic security can and should be passed to cognizant authorities for disposition under legal statutes.

Sharing early knowledge of China's AI research supports U.S. competitiveness.

China's STI-monitoring network operates less as an early-warning system and more as a siphon for transfer opportunities. By contrast, the U.S. government has long refused to share foreign intelligence with private enterprise, even where such knowledge would accrue to U.S. national advantage, in part because there is no fair way to apportion the benefit between companies, and in part because the United States could get by without doing it. That situation no longer obtains. China is a peer competitor in many advanced technologies and it may be time to abandon the hubris that puts U.S. companies at a disadvantage.

There is bipartisan will to meet China's technological challenge and support remedies.

As a practical matter, given the U.S. intelligence community's reluctance to build robust open source–based foreign technology detection and analysis programs,[235] the only economically viable way to build a risk and threat detection mechanism is to start "small" (one country), test methodologies, establish a reputation and user base, and thereby attract additional support to expand in increments to a full-blown operation that monitors all state and non-state actors' advanced AI developments. In today's milieu, China topics draw the most attention and hence are more likely to be funded.

China's development of general AI carries potential existential risk.

Students of catastrophic risk agree that even if the chance of "runaway AI" (or a meteor strike, or fatal pathogen) is low, the potential consequences may be grave enough to justify strong efforts at mitigation.[236] Similarly, if the risk of China—or anyone else—creating or stumbling upon an asymmetrical advantage in AI is small, *and* confidence in safeguards is high, the results of error, misplaced trust, and bad judgment may be serious enough to warrant scrutinizing developments in advance.

An additional consideration is the technical requirements of open-source exploitation. Tracking China's S&T development—in AI or any field—cannot be done without Chinese-language sources. To believe otherwise is to imbibe yet another China myth. Indeed, most of the "good" items are buried in what frustrated analysts call the "soft encryption" of China's logographic script. Working with the native language and orthography presents challenges, which, if solved, take one most of the way to a universal data processing capability. Adding other languages and countries is straightforward by comparison.

G. CREATING A DATA MODEL OF THE PROBLEM

Building a scalable indications and warnings (I&W) watchboard of foreign S&T developments begins by populating a relational database of entities, attributes, and connections—in essence, a software model of the target. The steps include defining search objects, identifying sources and arranging (continuous) access, compiling an initial dataset, expanding it, expressing different data types in a common format, conditioning data to eliminate errors and variants (entity resolution), purging extraneous data, building a user interface, "tuning" the system to flag data that meet thresholds, establishing routines for supervision and service, and building a customer base and alerts protocol.

This sequence is an abstraction. In real-world settings, the steps overlap. Best practices use a combination of humans and machines, one supporting the other, according to needs that develop in practice and change constantly. The notion that an operation like this can be executed from a priori assumptions on the basis of "data science" is a fantasy engaged in by people who never struggled with "dirty" data mined from strategic competitors who speak different languages and whose interests usually conflict with those building the database.

Taking this principle into account, this paper's authors compiled a list of source and method *types* based on prior success at eliciting Chinese-language data on advanced technologies[237] and AI in particular.[238] They include a dozen "general" types accessible to most linguist-researchers that were used to build the starter list of entities in this paper's Section 2 and another half dozen "advanced" types that are more resource intensive.[239]

Each source type is serviced by a combination of human and artificial agents. In addition, information from the following sources was added to the starter list:

- Chinese presenters at international and China-based "AGI" conferences from 2009 to present with affiliations and biographic data.
- Chinese scientists working in any of China's BI-AI institutes from internet searches and confirmed by email address domains.[240]
- metadata on authors, affiliations, and research from *limited* keyword-based queries of Chinese Academic Journals (CAJ).[241]

Data from these four sources—the starter list and three accretions—was entered into Excel spreadsheets formatted to populate a relational database of organizations and persons (chiefly scientists and science managers).[242] This does not make up an I&W watchboard. It allows (1) the *accumulation* of topical information, (2) access to conditioned (clean, searchable, reliable) data, and (3) building of *network diagrams* based on location, co-authorship, and citation patterns that support discovery and analysis.

There is a learning curve in database construction—a tension between the needs of the analyst and the software engineer, best adjudicated when the volume of data is sufficient to bring most issues to light (e.g., how many categories of what type are needed) but before the data grows to a point where major changes to the data structure are unfeasible. These issues are being sorted out in preparation for the next stage of data acquisition, which involves full-scale exploitation of some, but not all, of the source types identified, namely:

- expanded keyword searches of the CAJ (CNKI) corpus.
- internet and news searches of key elements already identified.
- website exploitation (scrapes) of entities surfaced in the initial data pulls.

While these tasks can be performed by human operators with requisite language skills, all can be augmented by automated systems to a greater or lesser degree. Accordingly, the authors defined and began implementing a monitoring tool specifically tailored to detect and highlight any mention of certain key technologies, persons, or organizations appearing in online news streams. This tool will serve as an alarm and triaging mechanism for human researchers to create analytic products and identify newly emerging trends as soon as they appear.

Managing the collected data presents its own problems. The authors extracted, transformed, and loaded (ETL) listed data sources into machine-readable structured datasets—a non-trivial task—and are designing and optimizing full-text search criteria to match keywords, persons, and organizations for use in automated searches against our full-text data collections. Finally, they are collating and curating matched results into analysis-friendly formats for tabular presentation or visualizations.

Automated support for data conditioning and entry is also a priority in this pilot project. For certain structured data, we have begun scraping and processing large volumes of academic publication metadata and academic and commercial conferences and expositions for reference in a single user interface. CSET researchers have also designed data-entry templates that facilitate manual curation of authors and organizations metadata meant to capture their key aspects. Populated templates go through a quality control process for entity resolution and system-wide accuracy.

H. WATCHBOARDING AND MONITORING

Subsection G (above) describes a database to support I&W on a particular problem—China's development of advanced AI. Creation of a database mirrors work done by the authors elsewhere on similar Asian-language datasets, so success was expected. A database, however, is not a "watchboard," which is uncharted territory.

While indicators exist for WMD and other technologies, there is no ready-made compilation for "general AI" because there is no consensus on the goal and no precedents to provide guidelines. Also "watchboarding" differs from "monitoring" in that the former implies progress toward a defined end, while the latter is largely "observation without expectation." The authors' present goal straddles the two: China's progress toward advanced AI by one of three paths is a "well-defined forecasting target" but hardly qualifies as "objectively and unambiguously evaluable."[243]

The aim at this stage, accordingly, is not to elicit quantitative measures of progress—although that is likely to change—but to determine who in China is taking what steps toward general AI, as shown by overt expressions and other common measures. The AI safety literature is replete with hypothetical risks and disaster scenarios, from which indicators can be inferred. There is also copious literature on AGI and its presumed requirements. Other indicators are generic and apply to any advancing technology enterprise.

The tasks come down to: (1) compiling indicators of progress toward general AI and (2) applying them in some systematic, sustainable fashion to the sources available. Task 1 is fulfilled by three lists of indicators and topical keywords generated from the authors' AI and AGI readings, including "indicators" of a general nature, potentially problematic research areas, and a bilingual keyword list (see Appendix).

The first two lists orient analysts toward areas that merit attention and serve as a basis for generating topical keywords. These keywords are proxies for the target and are run periodically against three data stores, namely, the project's relational database, large proprietary data holdings—in this case, 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[244]—and the internet itself via purpose-built crawlers and newswire databases. The output is local alerts. Associated data may or may not be added to the database depending on what is retrieved, since the goal is to *model* the target, not replicate it.

For select internet searches, we first compile a list of query terms designed to optimize the retrieval of relevant search results. We implement heuristics for result relevancy, and then we query via API calls to major internet search engines.[245] Finally, we deduplicate and aggregate relevant search results in a human-friendly format for further triaging of analysis-relevant results.

For news monitoring, we compile a list of query terms designed to optimize the retrieval of relevant news article mentions. We collect Google News RSS and LexisNexis feeds to build custom newswire databases for each topic. This allows us to plot the number of news mentions of topic keywords over time.

The three lists are fluid documents adjusted for their ability to elicit material and our evolving understanding of the target. We invite expert participation in their growth and development. We are also experimenting with algorithms that can nominate new search terms from context.

The aim is an optimized keyword list able to absorb sophisticated metrics, including genuine numerical benchmarks. To stay relevant in a shifting technological landscape, we propose using intermittent surveys of AI experts from the academic community to inform algorithms that weigh how information and data are displayed in the monitoring dashboard tool. Term weighting will help ensure significant articles and news developments are put front and center for human dashboard monitors. As a practical matter, bounds for the keyword list will fluctuate between 100 and 200 terms, including names of selected scientists, by analogy with the Swadesh lists used to isolate key concepts from open-ended vocabularies.[246]

4

The China AI Monitoring System in Operation

This pilot project to monitor China's progress toward advanced artificial intelligence will be fully operational when coupled with stable ownership and resourcing, expert benchmarking, and mature processing algorithms. Meanwhile, this section offers two examples—a success and a failure—of the authors' early efforts to monitor indications of advanced Chinese AI programs, along with a compilation of retrieved statements on "general AI's" trajectory from Chinese experts identified by the project.

I. CASE STUDIES IN MONITORING AND DISCOVERY

Xiamen University and AGI

Seth Baum, in his 2017 survey of worldwide AGI programs, de-listed a program at China's Xiamen University (厦门大学) run by AGI pioneer Hugo de Garis from his inventory.[247] The university did not reappear in Baum's updated survey a few years later,[248] but remained of interest because of the prior activity.

Accordingly, the authors ran searches in 2021, both in the aforementioned database and online, aided by knowledge of China's AI programs accumulated during the general study. The searches revealed that the university had stood up a reinvigorated "Institute of Artificial Intelligence" (人工智能研究院), with an unusually high number of tenure (track) positions among its "AI-related faculty."[249] Other datapoints hinted at the university's intent to work toward AGI:

- The AI institute researches machine vision, pattern recognition, and cognitive science, and has teams studying hybrid intelligence and cross-media AI—all AGI precursors. [250]

- The institute has cooperative research relationships with all four of China's "BATH" (Baidu, Alibaba, TenCent, Huawei) companies, each of which has AGI research programs.[251]

- AI institute professor Ji Rongrong (纪荣嵘) lamented the inability of AI to deal with the "catastrophic forgetting" problem. Solving it is the "key to AGI" (通用人工智能).[252]

- Xiamen University's AI institute and Tencent, a known AGI developer,[253] announced jointly that "multi-modal fusion," in which the institute specializes, will lead to AGI.[254]

- Ben Goertzel is a visiting professor at Xiamen University's Fujian Provincial Key Laboratory of Brain-like Intelligent Systems (福建省仿脑智能系统重点实验室).[255]

- A high-profile paper by Xiamen University researchers on "Bioinspired Nanofluidic Iontronics" describes brain-like signal processing for neuron-computer interfaces.[256]

- Dean of the College of Humanities Zhu Jing (朱菁) and colleagues provided slides at the 2020 China AGI Annual Conference on the definition of AGI.[257]

- Jiang Min (江敏), professor in the university's School of Informatics and a senior IEEE member, co-authors with Goertzel and focuses personally on software and AGI.[258]

Other indicators of an AGI program were found in academic publications and collaboration networks, retrieved from the CSET merged corpus.

FIGURE 1

Papers by Xiamen University researchers on AGI topics

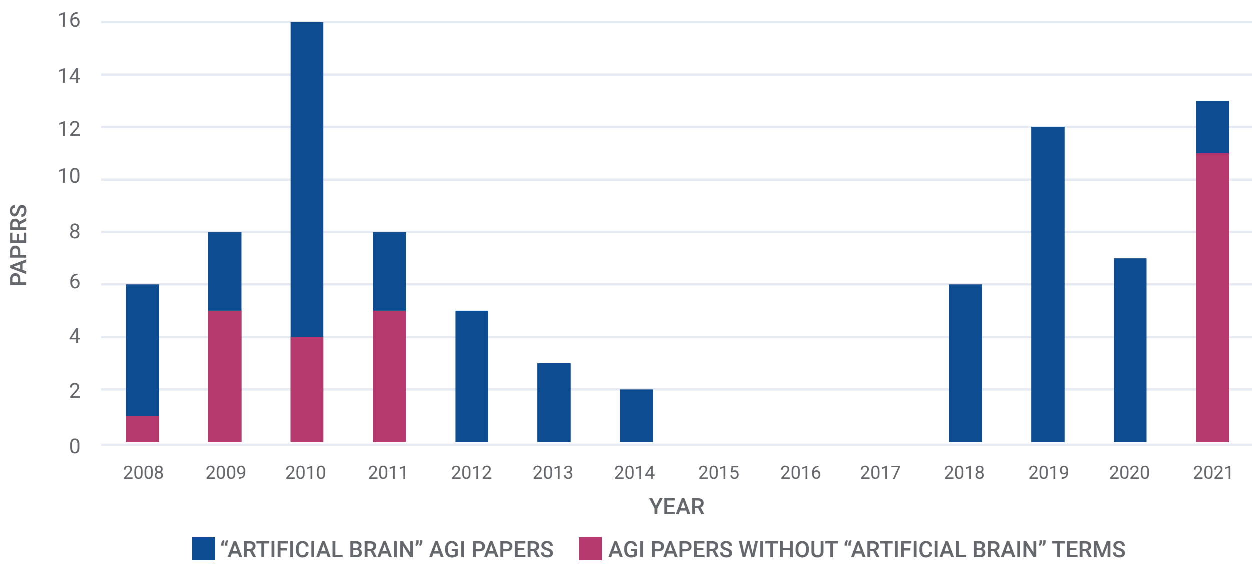

SOURCE: CSET MERGED CORPUS

Figure 1 shows the distribution by year of AGI-themed papers produced by Xiamen University affiliates starting in 2008, when de Garis' project began, through its termination in 2011 and on to mid-2021 (data cut-off). The blue areas are papers retrieved from CNKI and English-language databases based on a half-dozen terms aligned with de Garis' China-Brain Project; the red areas represent papers retrieved by searches on Chinese and English expressions for "AGI." In terms of scholarly output, the program's disappearance had a temporary impact only on (overt) AGI research at the university or among its affiliates.

TABLE 1

Papers co-authored by Xiamen University and major Chinese organizations with known AGI research programs

ORGANIZATION	PAPERS
Baidu	24
Alibaba	11
Tencent	93
Huawei	46
CASIA	16

SOURCE: CSET MERGED CORPUS

Table 1 shows Xiamen University co-author relationships with affiliates of China's big four AI companies and the Chinese Academy of Sciences' Institute of Automation, which all have or will have AGI research programs. The papers show the existence of collaboration networks only. Figure 2 is the same information expressed over years (data available to June 2021 only).

FIGURE 2

Xiamen University Collaboration with BATH + CASIA

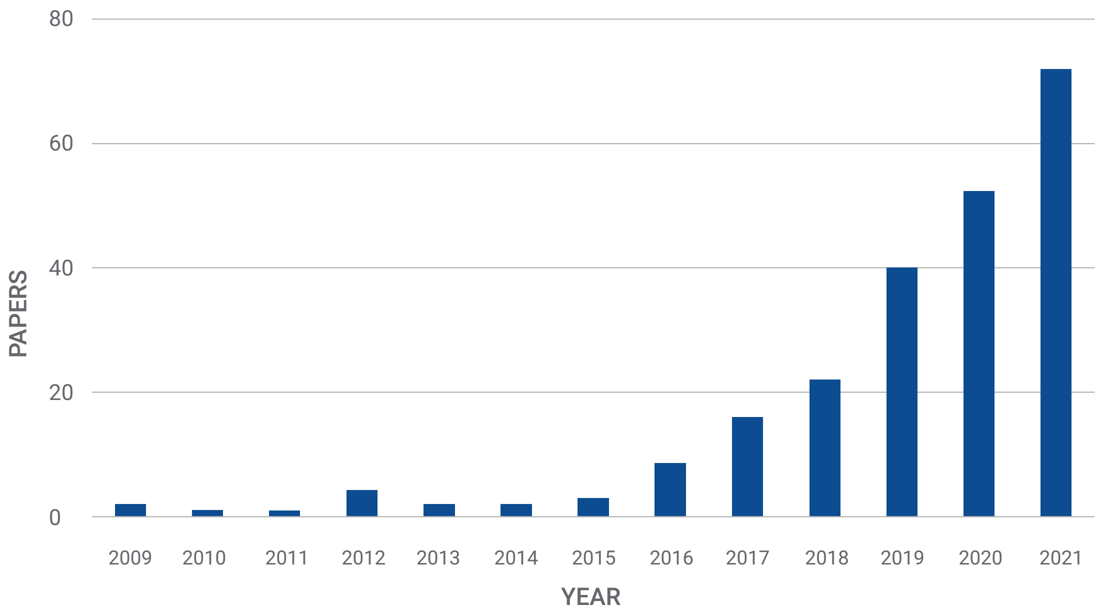

SOURCE: CSET MERGED CORPUS

Finally, we doubt that Professor de Garis failed to bequeath a legacy. His "China-Brain Project" involved 20 people, several still there,[259] building their "network of networks" vision of brain function.[260] Although we did not turn up a "hot" AGI research program at Xiamen University, this is likely a distinction without a difference, as there is ample evidence this provincial university is moving (back) toward AGI research, if it ever left it.

Wuhan's AI Research Complex

Whereas the authors were able to document the likely return of AGI research to Xiamen, an event of similar—or greater—import occurring in Wuhan went unnoticed until recently for lack of a robust monitoring protocol. The People's Republic of China has a history dating from the 1950s of shunting sensitive research to the hinterlands.[261] This may or may not apply to artificial intelligence, but habits persist, so when significant developments in AI happen in the provinces, outside the eastern complexes (Beijing, Shanghai), one should take notice.

For example, Table 2 was compiled from a corpus of Chinese academic journal articles in mid-2021 as a baseline for identifying key Chinese AI institutes. The table ranks PRC organizations by number of CNKI AI publications in 2020; the 2015 rankings are added for comparison.

TABLE 2

Organizations with top AI research output in CNKI journals

2020 RANK	2020 PAPERS	2015 RANK	2015 PAPERS	ORGANIZATION
1	1316	2	777	University of the Chinese Academy of Sciences (中国科学院大学)
2	1249	1	1037	**Wuhan University School of Information Management** (武汉大学信息管理学院)
3	1132	4	713	Tsinghua University Institute of Automation (清华大学自动化系)
4	928	6	616	Shanghai Jiao Tong University (上海交通大学)
5	889	18	502	Sichuan University College of Computer Science (四川大学计算机学院)
6	795	31	386	PKU School of Electronics Engineering and Computer Science (北京大学信息管理系)
7	770	80	221	Renmin University of China (中国人民大学)
8	770	5	680	Zhejiang University (浙江大学)
9	762	14	528	Tongji University (同济大学)
10	732	15	527	University of Shanghai for Science and Technology (上海理工大学)
11	718	46	311	Beijing Normal University (北京师范大学)
12	715	27	435	Nanjing University (南京大学)
13	701	3	748	NUAA (南京航天航空大学)
14	688	16	522	Jilin University (吉林大学)
15	686	12	531	Southeast University (东南大学)
16	658	10	567	Tianjin University (天津大学)
17	632	42	345	**Wuhan University of Technology** (武汉理工大学)
18	607	20	492	North China Electric Power University (华北电力大学)
19	605	23	469	Southwest Jiaotong University School of Electrical Engineering (西南交通大学电气工程学院)
20	595	25	459	**Huazhong University of Science and Technology** (华中科技大学)

SOURCE: CSET MERGED CORPUS

The authors, who have followed AI for years, were surprised by Wuhan University's School of Information Management's outsized contribution—which should have, but did not, surface in prior research—complemented by two other Wuhan institutes in the top 20 (Wuhan-based organizations are in bold).[262]

Wuhan's Huazhong University of Science and Technology (number 20) attracted our attention earlier when its HUST-Suzhou Institute of Brainsmatics received disproportionately large funding from China's National Natural Science Foundation compared to other recipients. Table 3 is a schedule of grants by NNSF for AGI-precursor projects, compiled by the authors in 2020 from incomplete data.[263]

TABLE 3

NNSF academic grants for AI-brain projects 2018–19

NNSF INDIVIDUAL GRANTS	2018		2019	
	GRANTS	FUNDING (USD)	GRANTS	FUNDING (USD)
BI-AI	31	2,524,531	19	2,002,811
Connectomics	18	3,959,650	11	1,398,454
Brain-computer Interfaces	20	1,655,108	19	2,202,753
Total (AI-Brain)	69	8,139,289	49	5,604,018
Average Grant Amount per Project	$117,961.00		$114,368.00	

SOURCE: www.letpub.com.cn/index.php?page=grant

The grants, small by U.S. standards, are typically matched by local resources or supplemented in other ways, but the main point is that two *other* NNSF grants were excluded from the tallies because of their unusual size:

> "Whole-brain mapping system based on morphology and omics spatial information" (基于形态与组学空间信息的细胞分型全脑测绘系统) for USD $10,224,342 awarded to Luo Qingming (骆清铭),[264] and

> "High-resolution optical imaging and visualization of brain connections in brain spatial information" (脑空间信息中脑连接的高分辨光学成像与可视化研究) for USD $2,800,821 to Li Pengcheng (李鹏程).

Professors Luo and Li were both at HUST in Wuhan. Luo, founder of its Brainsmatics Institute, was also dean of the Wuhan Optoelectronics National Research Center of HUSTechnology,[265] engaged mainly in photonics research including "multi-modal molecular imaging" in support of BI-AI and connectomics.[266] The import of this funding anomaly did not register with us then, either, and we treated it as an outlier.

In retrospect, other signs of Wuhan's emergence as a potential AGI center were there had we been looking. A professor at Huazhong Normal University (华中师范大学), Liu Kai (刘凯),[267] has sponsored China AGI Annual Conferences (中国通用人工智能年会) since 2016,[268] a result of Liu's collaboration with U.S. AGI expert Pei Wang.[269] The conferences supplement the Artificial General Intelligence Society's series, which has run worldwide since 2008 to showcase local developments, and a disproportionate number of its attendees are from Wuhan institutes.[270]

The city made AI news in May 2021, which we also missed, when the Wuhan Artificial Intelligence Computing Center (武汉人工智能计算中心) began operation, billed as the start of "a new generation of AI innovation."[271] That month, the CAS Institute of Automation (a declared AGI developer), Huawei, and the city's East Lake Hi-tech Development Zone (东湖高新区) announced plans to build a Wuhan General Artificial Intelligence Platform (武汉通用人工智能平台), create a "perception-cognition-decision whole chain ecology," and an open platform for "autonomous and controllable general AI" (自主可控通用人工智能).[272]

This development was impossible to ignore. In July 2021, CASIA announced a "multi-modal general artificial intelligence" project called "Zidong Taichu" (紫东太初) based on Huawei's "Shengteng" (昇腾) platform, aimed at integrating image, text, and audio data to "replicate human semantic processing." Computing support is via the Wuhan AI Computing Center.[273]

Wuhan's role as a key node (核心节点) and next step (下一步) in AI development was made explicit, for anyone listening, on December 20, 2021 by CASIA's Wang Jinqiao (王金桥), who acknowledged that OpenAI's GPT-3 (subsection A above), from which Zidong Taichu draws its inspiration, "opens a new beginning for artificial intelligence from dedicated intelligence to general intelligence"[274]—the whole point of our discovery project.

Could we have forecast Wuhan's emergence as an advanced AI center earlier, with the proper data, protocols, and incentive? What other potentially transformative developments have we and other observers been missing?

J. CHINESE EXPERT VIEWS ON "GENERAL AI"

This subsection assesses China's willingness to *accept* advanced AI, based on expert views discovered while executing this project. Chinese AI scientists, like their counterparts worldwide, are aware of the controversies surrounding the definition, attainability, and risks of "AGI." Their views on these matters will impact practical outcomes. First some caveats:

The authors acknowledge the larger issue of whether acceptance of advanced AI per se by China's governing class, S&T policymakers, AI scientists, and informed public even matters. An argument can be made that technology follows its own dynamic—what can be done *is* done, driven by pressure to accept innovations that offer a higher "utility function" no matter what its long-term outcome may be. "Choice" on a macroscale may be as illusory as the alleged "free will" Chinese AI scientists (E.9 above) propose to investigate.

Technical innovation may be *delayed* by government, for good reasons (environmental degradation, other existential threats), bad reasons (elite paranoia, societal collapse), or in the case of the United States, for no reason (hubris, complacency). In an authoritarian country especially, the speed at which events happen is determined to a greater degree by the attitude of its governing body. We will argue that China's political elite not only support advanced AI but see it as in their class interest to continue doing so.

Finally, acceptance itself becomes a self-fulfilling prophecy. Declaring a goal of world domination by 2030 may spur China's AI scientists to act in ways that create the desired end, for example, by ignoring safety concerns that hobble other nations or simply by believing in a given outcome. As Goertzel put it, China "has no 'chip on its shoulder' about AI or AGI—it has the same status as any other advanced technology."[275] So Zhu Songchun, Nikos Logothetis, and others gravitate there and help drive China's success.

What do Chinese scientists think of general AI's prospect and how comfortable are they with it? The authors cite two surveys, bearing in mind that the focus is "advanced AI" with general capabilities, not the human level AGI addressed by the surveys, which may (or may not) differ.

A paper by Grace et al. in 2018 reports a survey of international researchers who published at two top machine learning conferences in 2015.[276] The forecasts per se (guesswork in any case) are less interesting than the differential findings for Asian and, especially, Chinese respondents, who are more optimistic about AGI's emergence:

> "(R)espondents from different regions had striking differences in HLMI [human-level machine intelligence] predictions. Fig. 3 shows an aggregate prediction for HLMI of 30 years for Asian respondents and 74 years for North

Americans. Fig. B.1 displays a similar gap between the two countries with the most respondents in the survey: China (median 28 years) and USA (median 76 years)."[277]

Mindful that beliefs matter, the present authors surveyed two groups of Chinese AI researchers in 2020 on issues related primarily to BI-AI. A majority (74 percent) stated that brain-inspired AI will lead to AGI (83 percent of the BI-AI specialists). AI generalists were split 50/50 on whether AGI would be achieved (through any means) in 5–10 years or more than 10 years, while BI-AI specialists were more aligned that AGI is more than 10 years away.[278]

While this data is indeterminate, the probability that Chinese scientists *imagine* and *contemplate* general AI's appearance at a relatively early date seems hard to refute. Some of China's leading AI scientists lend credence to the argument that China's research establishment both welcomes and is working toward advanced (general) AI:

Dong Le (董乐), deputy director, Beijing Institute for General Artificial Intelligence, at the BEYOND International Technology Innovation Expo in Macao, December 2021.

> The trend of AI is from perceptual to cognitive intelligence. We are following a pathway in the "evolution of general artificial intelligence" with the potential to "become the most reliable and important aid in the human-machine merger of the future."[279]

Gao Wen (高文), dean of Peking University's School of Information Science and Technology and director of Pengcheng Lab (C.6 above).

> AI is at a critical stage in the development of a new generation of artificial intelligence to strong artificial intelligence. By 2030, China's artificial intelligence will reach the world's leading level.[280]

Huang Tiejun (黄铁军), vice dean of Peking University's Institute for Artificial Intelligence and an outspoken advocate of ASI.

> Artificial superintelligence is the future—an "evolutionary trend." Technological innovation and development should not be restricted by the limitations of human beings. We can treat the development of artificial intelligence with a more open mind.[281]

Ming Dong (明东), Tianjin University professor and director of its Brain Science and Brain-like Research Center (E.4 above).

> "The future is not about replacing human beings with artificial intelligence, but making AI a part of human beings through interconnection and interoperability. A blend of human and computer without barriers is the inevitable end of the future."[282]

Pan Gang (潘纲), Zhejiang University professor of computer science, AI, computer vision, and BCI expert.

> BCI will support "a new type of intelligence—brain-computer hybrid intelligence" (脑机混合智能) able to effect comprehensive enhancement of one's perceptual, cognitive, and behavioral ability, a "mixture of biological and machine brain."[283]

Shi Luping (施路平), director of the Center for Brain-inspired Computing Research, Tsinghua University and head of the Tianjic team.

> "Nanodevices have enabled us to develop electronic devices such as neurons and synapses at the level of human brain energy consumption, so now is the best time to develop general artificial intelligence."[284]

Tan Tieniu (谭铁牛), former deputy director of the Chinese Academy of Sciences and deputy chief of the PRC's liaison office in Hong Kong.

> "How to make the leap from narrow artificial intelligence to general artificial intelligence is the inevitable trend in the development of the next generation of artificial intelligence. It is also a major challenge in the field of research and application."[285]

Xu Bo (徐波), director of CAS Institute of Automation and chair of China's Next Generation Artificial Intelligence Strategic Advisory Committee.

> "General artificial intelligence has always been a dream in the technology world ... We believe autonomous evolution [a topic of the institute's research] is a bridge from weak artificial intelligence to general artificial intelligence."[286]

Yao Dezhong (尧德中), director of the MOE Lab for Neuroinformation (神经信息教育部重点实验室) in Chengdu on development of a "digital twin brain (数字孪生脑) model."[287]

"It is believed that the future digital twin brain can become a general intelligent system in a certain sense by imitating the working principle of the human nervous system, bringing AI to a new level. Thinking and decision-making are close to the human brain."[288]

Zhu Songchun (朱松纯), director of BIGAI, former UCLA computer vision expert, proponent of the "small data, big task" approach to AI development.

"The goal of general AI research is to create general agents with autonomous perception, cognition, decision-making, learning, execution, and social collaboration capabilities that conform to human emotions, ethics, and moral concepts."[289]

Zeng Yi (曾毅, C.9 above), prominent AI safety advocate, speaker at the Future of Life Institute's "Beneficial AGI 2019" conference, and head of a BAAI group that developed the "Beijing AI Principles,"[290] whom one would expect to be against general AI, is on record *backing it.*

"The development of narrow (专用) AI does not completely avoid risk, because the system is likely to encounter unexpected scenarios in its application. Having a certain general ability may improve the robustness and adaptiveness of an intelligent system."[291]

Support from the scientific establishment does not guarantee acceptance by society. Outside of China, AI (not to mention AGI or ASI) conjures up dystopian scenarios among many informed persons (Hawking, Musk, Bostrom, and Gates)—concerns that are well-founded. While the authors share these concerns, overcaution can mean ceding to one's competitor the "first mover advantage" (先发优势) China demonstrably seeks,[292] which in AI may be difficult or impossible to close. So it is worth contemplating how receptive China as a whole is to an expanded role of AI in human life.

While the authors know of no popular surveys, China's acceptance of general AI is the topic of two recent books. Song Bing (宋冰), vice president of the Berggruen Institute and director of its PKU-based China Center (北京大学博古睿研究中心) published an edited book by 17 Chinese philosophers and scientists titled 《智能与智慧: 人工智能遇见中国哲学家》 (*Intelligence and Wisdom - Artificial Intelligence Meets Chinese Philosophers*),[293] billed as "the first systematic endeavor by prominent Chinese philosophers . . . to address challenges and opportunities posed by frontier technologies such as artificial intelligence and robotics."[294]

Song, a former Goldman Sachs executive and capital markets lawyer presumably not given to flights of fantasy, summarized the book's content in interviews. As she sees it, "Western media and intellectual elites generally have more vigilance and fear of cutting-edge technology," and of AI especially, than Chinese, which she attributes to Western human-centrism—in contrast to the "non-anthropocentrism" (非人类中心主义) of China. So "the emergence of an existence stronger than human is not a problem."[295] "What is wrong coexisting with machine superintelligence (超级机器智能)?"[296]

Song is correct ascribing to China a preference for holistic thought that denies humanity's privileged status.[297] These contrasting *tendencies* of Eastern and Western thought have been known for decades and observed in careful psychological studies.[298]

The second book, by Fudan University philosopher Xu Yingjin (徐英瑾) titled 《人工智慧哲學十五講》 (*Fifteen Lectures on the Philosophy of Artificial Intelligence*),[299] likewise comes down unequivocally in favor of strong AI:

> "Does the ethical benefit of researching general AI outweigh the cons, or does the harm outweigh the benefits? My answer is 'the pros outweigh the cons' (利大于弊)."[300]

The ultimate gauge of AI's acceptance is the attitude of China's ruling class, as measured by institutional support, public declarations, and by the relationship of artificial intelligence to the needs of political dictatorship. The first of these, state support, is documented in a 2022 book edited by the present authors,[301] especially its chapter on "State plans, research, and funding."[302] Support for advanced AI is also explicit in China's 2017 foundational AI document, the "New Generation AI Development Plan,"[303] which lays the groundwork for "general artificial intelligence," (通用人工智能), while calling for a merger (混合) of AI and human intelligence.[304]

Endorsements by China's top political leaders are easily found. On May 30, 2016, President Xi Jinping in a speech to a MOST-sponsored assembly titled "Striving to Build a World S&T Superpower" made the following reference to one of the three advanced AI pillars in China:

> "Connectomics is at the scientific forefront for understanding brain function and further exploring the nature of consciousness. Exploration in this area not only has important scientific significance, but also has a guiding role in the prevention and treatment of brain disease and the development of intelligent technology."[305]

Xi's report to the 19th Party Congress in October 2017 highlighted progress in AI and quantum science.[306] Since Xi assumed power in 2012, only three of 70-odd Politburo study sessions have been on specific technologies—AI (2018), blockchain (2019), and quantum computing (2020).[307] CAS Institute of Automation's website lists 13 instances from May 2018 to November 2020 where Xi publicly expressed support for AI.[308] For example:

> "Driven by new theories and technologies such as mobile internet, big data, supercomputing, sensor networks, and brain science, artificial intelligence has accelerated its development, causing the emergence of deep learning, cross-discipline integration, human-machine collaboration, group intelligence, and autonomous control."[309]

In Xi's estimate, AI has the effect of a "lead goose" (头雁), i.e., a spillover function driving the current S&T "revolution" and "industrial transformation (产业变革)."[310]

Are there other reasons for Chinese leaders' embrace of AI? It is an open question whether advances in information technology benefit freedom (cryptocurrency, blockchain, virtual private networks) or dictatorship (censorship, predictive policing, ubiquitous surveillance). AI-enabled influence operations are also a two-way street. What seems certain is that whichever "side" has a technical advantage is primed to prevail. As the struggle for information dominance becomes automated, the potential for AI surprise will increase, raising the stakes.

The Chinese government's use of AI as an instrument of oppression is documented by Dahlia Peterson and Maya Wang, among others, and requires no elaboration here.[311] It is also a topic of study at CSET, potentially as a follow-up database exercise.[312] The point is that China's Communist Party leaders' acceptance of advanced AI, beyond enhancing national power, may be motivated by a *need* to own it just to hold the line.

Recommendations

China, through complementary approaches and state backing, is vigorously pursuing advanced general AI, understood as broadly capable software that can function autonomously (creatively) in novel environments. Whether this research leads to human-level AI—however defined—is irrelevant, as the impact of these advances on China's ability to survive as an authoritarian state while projecting its will globally will be comparable on many levels to scenarios associated with high-level machine intelligence. Indeed, if China's programs in BCI and connectomics succeed, the distinction between human and machine intelligence may disappear.

What might frustrate China's goals? We acknowledge, with China's S&T community,[313] that basic science has not been the country's strong suit, but find little comfort given (1) a paradigm shift in the quality and quantity of China-authored scientific papers, (2) China's success in facilitating the return, importation, and cooperation of foreign-trained scientists, and (3) China's storied ability to operationalize scientific discovery in advance of those making the discoveries, which in the end is all that matters.[314] These circumstances take on added urgency in light of AI's function as an enabler of most other scientific endeavors.

China's intent to lead the world in AI by 2030 and achieve a "first mover advantage" cannot be dismissed as rhetoric but is a serious prospect that will have practical consequences. Of course, China's future capabilities *relative to other nations* is a key element this study does not address, although it is obvious that projections of any sort depend on the availability of data and one's ability to derive from it *indications* that can be assessed.

This paper outlines a monitoring project meant to kickstart an I&W regime. The need is not limited to China or AI but applies to foreign S&T in general. Sadly, such an open-source collection-and-analysis system does not exist anywhere among the liberal democracies—neither inside nor outside the U.S. intelligence community—an issue the authors and their colleagues have lamented in books, papers, lectures, congressional briefs, and other public and non-public outreach.[315]

Although these appeals have elicited principled agreement, the dialog invariably stops at how and where a capability can be stood up. No intelligence agency is suited to own it; there is no will at present to create an entirely new organization, let alone fund it; and "joint task forces" have half-lives measured in months. Never mind that China has been running STI operations against U.S. technology for six decades with a cadre of one hundred thousand and budget that, until recently, surpassed the state's own R&D expenditure.[316]

A solution that *could* take hold immediately is a modest expansion of this pilot project (a "minimum viable product"[317]) made up of the following:

- Two level 3+ Chinese linguist data collectors, primarily for internet exploitation
- A level 2+ linguist with topical expertise, for commercial data holdings
- A level 3 linguist to perform quality control and local database entry[318]
- A non-linguist database manager with "proficient" level IT skills[319]
- A software engineer/data scientist with "expert" programming skills
- Two level 3 linguist-analysts with developed writing and briefing skills[320]
- A product editor and graphics design specialist
- One project manager able to engage s/w engineers and collectors/analysts
- A target savvy outreach person able to interact with U.S. and global experts

An operation of this scale is sufficient to produce "indications and warnings" on a timely basis for China AI *or any other priority discipline*, while establishing the viability of the construct in general—i.e., the project is inherently scalable, or it can function as is indefinitely.

Although our preference is to situate the project in an open or sensitive-but-unclassified setting, the authors are willing, with our Center's blessing, to transfer methods and expertise to any non-commercial U.S. or allied organization able to step up to this task.

A final recommendation, better viewed as a caution, pertains to China's pursuit of a "first mover advantage." We have highlighted the challenge that success here will present to China's global competitors but have not sufficiently emphasized that China's current "*second* mover advantage" is cause for concern by itself, given China's legendary expertise at spotting foreign innovations and highly effective infrastructure for expediting their transfer. Whatever label one attaches to the reality, the ability to monitor China's technological development and convey accurate judgments to policymakers will determine whether liberal democracy survives these challenges.

Appendix: Indicators and Keyword Lists

The following is a first attempt at composite lists of advanced AI indicators (Lists 1 and 2) and keywords (List 3). Work is ongoing to determine the optimum between too wide and narrow.

List 1: Potential indicators—general

activities, statements of dedicated organizations
advances in chips/neuromorphic computing
algorithmic advances
AI safety research (or announcements)
applications of AI in new and significant areas
"black" AI programs
"Chinese room" (Searle) or Turing Test claims/successes
collaboration with advanced foreign AI companies (esp. those doing AGI)
concentrations (movements) of known actors
"disappearance" of significant researchers or lines of research
domestic or international AGI seminars
expansion of physical plant
expressions of intent[321]
hacks of tech firms (IP exfiltration)
hacks of data holdings (e.g., OPM or Facebook for training data)
hardware advances (in general)
importation of specialists (foreign or "returnees")
improved neuroimaging and connectomics sequencing
increased budgets (government and private sector)
increases in compute speed (esp. through quantum or nanotech)
increases in/concentrations of compute resources (including virtual networks)
investment decisions
new data-reshoring requirements (for access to training data)
procurement red flags (e.g., AI/FR software by, e.g., MPS, MSS)
published statements about an "AI arms race"
related advances in neuroscience and cognitive science
reward hacking[322]
"seven sons" and SASTIND university related research

"singleton" (霸) discourse

spikes in articles defining "intelligence"

uptick in non-human primate production and research

List 2: Potential indicators—research

accountability

advanced virtual reality

AI safety and control

AI "orthogonal" valuation

attention to the control problem

BCI chip implants, wearable BCI devices

building artificial environments for learning

creating or stealing an "AGI roadmap"

discreet modeling of human psychological development

embodiment and motivation research

explainability problem (XAI)

human chimera research

large integrated knowledge bases (architecture isn't enough)

ML aimed at reusing existing skills (instead of learning completely new skills)

ML for acquiring general skills vs. discreet (complex) problem solving

quantum computing (to speed up ML)

semantic understanding of language

staged (incremental) learning curricula

tailored brain-inspired AI research

whole brain emulation ("mind uploading")

List 3: Bilingual keywords

abstract event representation	抽象事件表示
active learning	主动学习
additive learning	加性学习
adversarial networks	对抗网络
affective computing	情感计算
AGI	
AI alignment	AI对齐
analogical reasoning	类比推理
anomaly detection	异常检测

architecture search	建筑搜索
artificial emotion	人工情感
artificial general intelligence	通用人工智能
artificial pedagogy	人工教学法
artificial psychology	人工心理学
artificial scientist test	(人工科学家测试)
ASI	
associative memory	联想记忆
attention modeling	注意建模
attentional memory	注意记忆
autobiographical memory	自传记忆
autonomous knowledge acquisition	自主知识获取
autonomous learning	自主学习
autonomous reasoning	自主推理
awareness	意识
blue brain project	蓝色大脑
brain in a vat/jar	缸中之脑
brain simulation	脑部模拟
brain transplant	脑部移植
brain-computer chip	脑机芯片
cascade/cascading	级联
category theory	范畴论
causal reasoning/inference	因果推理
cognitive architecture	认知架构
cognitive computing	认知计算
cognitive enhancement/nootropic	认知增强
common sense	常识
complete AI	完全 AI
computational creativity	计算创造力
computing overhang	计算悬垂
concept formation	概念形成
conceptual binding	概念绑定
conscious perception	有意识的知觉
consciousness	意识
consciousness and memory model	意识和记忆模型
consciousness uploading	意识上传
continual/continuous learning	持续学习

control problem	控制问题
cross-domain generalization	(跨域概括)
cumulative learning	累积学习
cyborg	赛博格
declarative knowledge	陈述性知识
declarative memory	陈述性记忆
discontinuous leap	不连续的飞跃 / 跃进
effective BCI	有效BCI
efficient concept formation	有效的概念形成
efficient learning	高效学习
embodiment	体现
episode retrieval	情节检索
episodic memory	情景记忆 / 场景记忆
evolutionary learning	进化学习
executive control	执行控制
feature binding	特征绑定 / 捆绑
flexible/elastic memory	弹性记忆
formal neural networks	形式神经网络
general-purpose system	通用系统
global workspace	全局工作空间
gradual neuron replacement	神经元逐步替换
hard takeoff	硬起飞
hierarchical decomposition	层次分解
higher order cognition	高阶认知
high-level machine intelligence	高级机器智能
HLAI	
HLMI	
human-level AI	类人 AI / 人类水平 AI
hypothetical reasoning	假设推理
imagination	想像力
incremental / progressive learning	渐进式学习
inductive inference / reasoning	归纳推理
integrated AI	集成 AI
integrative designs	一体化设计
intelligence explosion	知识爆炸
intelligence science	智能科学
intentional memory	有意记忆
intentionality	意向性

intuitive understanding	直觉的理解
knowledge representation	知识表示
knowledge transfer	知识传输 / 转移
lifelong learning	终身学习
long-term knowledge	长知识
long-term potentiation ML	长期 / 长时程增强/增益 ML
LTP machine learning	LTP机器学习
machine ethics	机器伦理
machine learning neurofidelity	(机器学习神经保真度)
mental representation	心理表征
metalearning	元学习
MicroPSI	
mind model(ing)	心智模型
mind uploading	心灵上传
mirror/mirroring	镜像
modeling of reasoning	推理建模
motive selection	动机选择
multi-objective reinforcement learning	多目标强化学习
natural language understanding	自然语言理解
neuroinformatics	神经信息学
neuron replacement	神经元替换
neurosymbolic representation	神经符号表示
Non-Axiomatic Reasoning System	非公理推理系统
non-linear dynamical systems	非线性动力系统
object permanence	对象永久性
object/scene vision	目标 / 场景视觉
one-shot learning	一次性式学习
OpenCog	
OpenPSI	
parameter	参数
pattern reconstruction	模式重构
perceptual consistency	知觉一致性
perverse instantiation	反向实例化
position invariance	位置不变
post-biological	后生物
predictive coding	预测编码
predictive modeling	预测模型

prisoner's dilemma	囚徒困境
procedural memory	程序记忆
procedure learning	程序学习
psyche + AI	心灵 + 人工智能
recursive self-improvement	递归自我完善
reinterpretation	重新诠释
relative strength index (RSI)	相对强度指数
reward replacement	奖励替换
seed AI	种子AI
self-reflecting	自我反省
self-awareness	自我意识
selfhood	个性 / 自我
self-improving	自我完善
semantic memory	语义记忆
semantic vision	语义视觉
sense making	意义建构
sensory memory	感觉记忆
shape sorting	形状分类
simulated environment	模拟环境
singularity	奇点
social intelligence	社会智慧
sparse representation	稀疏表示
spike timing dependent plasticity	脉冲 / 尖峰时序依赖可塑性
spiking neural network	脉冲神经网络
strong AI	强 AI / 强人工智能
superintelligence	超智能
technological singularity	技术奇点
thinking machine	思维机器
transfer learning	迁移学习
transhuman	超人的 / 超人类
transhuman(ism)	超人类(主义)
uncertainty estimation	不确定性估计
unified knowledge	统一知识
variable binding	变量绑定
whole brain emulation	全脑仿真

Endnotes

1. PRC State Council, 国务院关于印发新一代人工智能发展规划的通知 ("The New Generation AI Development Plan") SC (35), 2017.
2. Kai-Fu Lee, *AI Superpowers: China, Silicon Valley, and the New World Order*, Harper Business, 2018.
3. William Hannas and Huey-Meei Chang, eds., *Chinese Power and Artificial Intelligence: Perspectives and Challenges*, Routledge, 2022. See especially chapters 5 and 7–9.
4. The People's Liberation Army has no shortage of thinking on advanced AI topics, including general artificial intelligence. As of 2018, military academics at the PLA's Academy of Military Sciences had called for more extensive research into the use of AGI (强人工智能) and ASI (超人工智能) for intelligentized weapons development. See Li Dewei (李立伟) and Zhu Lianhong (朱连宏), "Considerations on Designing the Smart Equipment System (对展开装备体系设计的初步思考)," *China Military Science* 2018, no. 1, p. 118-119. Their more famous comrade and CAE academician Li Deyi (李德毅) also addressed the subject of general artificial intelligence. See Li Deyi (李德毅), "Ten Questions and Answers for the New Generation of Artificial Intelligence (新一代人工智能十问十答)," *CAAI Transactions on Intelligent Systems* 2021, Vol. 16, No. 5, p. 828-833; John Chen, personal communication, April 13, 2022.
5. Ngor Luong and Ryan Fedasiuk, "State Plans, Research, and Funding," in Hannas and Chang, *Chinese Power and Artificial Intelligence*; Karen Sutter and Zachary Arnold, "China's AI Companies: Hybrid Players," in Hannas and Chang, *Chinese Power and Artificial Intelligence*.
6. Hannas and Chang, "China's Access to Foreign AI Technology," (Center for Security and Emerging Technology, September 2019), https://cset.georgetown.edu/publication/chinas-access-to-foreign-ai-technology/; Hannas and Chang, "China's 'Artificial' Intelligence," in *China's Quest for Foreign Technology: Beyond Espionage*, Hannas and Didi Kirsten Tatlow, eds., Routledge, 2021; Hannas and Chang, "Foreign Support, Alliances, and Technology Transfer," in Hannas and Chang, *Chinese Power and Artificial Intelligence*.
7. Ryan Fedasiuk and Emily Weinstein, "AI in the Chinese Military," in Hannas and Chang, Chinese Power and Artificial Intelligence.
8. "Name rectification" (正名) is a Confucian injunction to avoid chaos by matching names with reality.
9. https://en.wikipedia.org/wiki/Intelligence.
10. Shane Legg and Marcus Hutter, "A Collection of Definitions of Intelligence," in *Advances in Artificial General Intelligence*, Ben Goertzel, Pei Wang, and Stan Franklin, eds., IOS Press, 2007.
11. Attributed to Charles Sparman in 1904.
12. E.g., creativity is regarded as an expression of intelligence and has been called the "ultimate moonshot" for AI (https://www.ibm.com/watson/advantage-reports/future-of-artificial-intelligence/ai-creativity.html). But its correlation with IQ is weak, disappearing altogether after about IQ 120 (Scott Findlay and Charles Lumsden, *The Creative Mind*, London: Academic Press, 1988. p. 36).
13. "The classical psychological measure of intelligence is the 'g-factor,' although this is quite controversial, and many psychologists doubt that any available IQ test really measures human intelligence in a general way." Ben Goertzel and Cassio Pennachin, "Contemporary Approaches to Artificial General Intelligence," in Artificial General Intelligence, Goertzel and Pennachin, eds., Springer, Berlin, 2007, p. 6.
14. Pei Wang, "The Logic of Intelligence," in Goertzel and Pennachin, *Artificial General Intelligence*, p. 32, 35. Overreach by AI's early pioneers led to "AI winters" of disillusionment and lost funding.
15. The authors recently briefed a U.S. government AI working group, whose first question was, "How do *you* define AI?"

16. Goertzel and Wang each have functioning AGI research programs; Yudkowsky's program is no longer active.
17. Wang, 2007, p. 31.
18. Ben Goertzel, "Artificial General Intelligence," *Scholarpedia*, 10 (11): 31847, 2015, http://www.scholarpedia.org/article/Artificial_General_Intelligence.
19. Eliezer Yudkowsky in Robin Hansen and Eliezer Yudkowsky, *The Hanson-Yudkowsky AI-FOOM Debate*, Machine Intelligence Research Institute, 2013, p. 346.
20. https://en.wikipedia.org/wiki/Artificial_general_intelligence; Goertzel, "Artificial General Intelligence;" https://searchenterpriseai.techtarget.com/definition/artificial-general-intelligence-AGI.
21. Alan Turing—The computer must fool a human into believing s/he is conversing with another human; Stephen Wozniac—The computer demonstrates an ability to make coffee in a stranger's home; Ben Goertzel—The computer passes enough university courses to obtain a college degree.
22. Muehlhauser and Bostrom's characterization of the imputed link between intelligence and morality is priceless: "It is an anthropomorphism that does not even hold true for humans." Luke Muehlhauser and Nick Bostrom, "Why We Need Friendly AI," *Think*, Spring 2014, p. 4.
23. Peter Voss, "Essentials of General Intelligence: the Direct Path to Artificial General Intelligence," in Goertzel and Pennachin, *Artificial General Intelligence*, 2007, p. 142.
24. Erik J. Larson, *The Myth of Artificial Intelligence: Why Computers Can't Think the Way We Do*, Harvard University Press, Cambridge, 2021.
25. AI pioneer Kai-Fu Lee reached the same conclusion: "I would suggest that we stop using AGI as the ultimate test of AI." Kai-Fu Lee, "Why Computers Don't Need to Match Human Intelligence," *Wired*, December 16, 2021, https://www.wired.com/story/deep-learning-versus-human-intelligence/.
26. E.g., Yang Zuomin's (杨作民) formulation of four stages of AI development: "weak AI" based on mathematics > "strong AI" based on algorithms and perception (感知) > "general AI" (通用人工智能) based on cognition (认知) > superintelligence. "人工智能技术在智能武器装备的研究与应用" (Research and Application of Artificial Intelligence Technology in Intelligent Military Hardware), 人工智能与机器人研究 (*Artificial Intelligence and Robotics Research*), 10.3, 2021.
27. China's major state-supported effort to achieve advanced (general) AI, whose director favors non-mainstream ("small data") approaches (see below).
28. OpenAI's third-generation "Generative Pre-trained Transformer" is a machine learning model trained on internet data to "perform a wide variety of natural language tasks." https://openai.com/api/.
29. Parts of Subsection B are mirrored in the authors' "Future Paths and Challenges," in Hannas and Chang, *Chinese Power and Artificial Intelligence: Perspectives and Challenges*.
30. Shane and Hutter, synthesized from 70+ definitions by philosophers, psychologists, and AI researchers, op cit.
31. Wang, 2007, p. 33; Also, 王培 (Wang Pei), "通用人工智能和非公里推理系统" (General Artificial Intelligence and NARS), summer course at Peking University, July 12–August 13, 2018, https://cis.temple.edu/~pwang/AGI-NARS/2018-PKU/AGI-NARS-2018.pptx.
32. Luke Muehlhauser, "Facing the Intelligence Explosion," Machine Intelligence Research Institute, 2013, p 43.
33. The discussion is limited to machine-based cognition and excludes other hypothesized ways to artificially elevate intelligence, e.g., drug-induced enhancement, swarm intelligence, gene-editing, and embryo selection. For the last, see Carl Shulman and Nick Bostrom, "Embryo Selection for Cognitive Enhancement: Curiosity or Game-changer?" *Global Policy*, 5.1, 2014, p. 85-92.
34. Goertzel and Pennachin, 2007, p. 22.
35. Hannas and Chang, "Future Paths and Challenges," in Hannas and Chang, *Chinese Power and Artificial Intelligence: Perspectives and Challenges*.
36. This formulation corresponds to Pennachin and Goertzel's second and third approaches, namely, those guided by the human brain and mind, with emphasis on the one or the other. Goertzel and Pennachin, 2007, p. 22.
37. Huang Tiejun (黄铁军), "Imitating the Brain with Neurocomputer: A 'New' Way towards Artificial General Intelligence," *International Journal of Automation and Computing* 14.5, October 2017.

38. Otherwise known as "whole brain emulation" (WBE, 全脑仿真) or "mind upload." This corresponds roughly to Pennachin and Goertzel's first approach to AGI: "that draw their primary structures and dynamics from an attempt to model biological brains." Goertzel and Pennachin, 2007, p. 22.
39. Huang, 2017.
40. Carl Shulman and Nick Bostrom, "How Hard Is Artificial Intelligence? Evolutionary Arguments and Selection Effects," *Journal of Consciousness Studies* 19.7-8, 2012.
41. This equates to Pennachin and Goertzel's fourth type of AGI project "that have drawn very little on known science about human intelligence in any regard." 2007, p. 22.
42. An "AGI may be built on computer science principles and have little or no resemblance to the human psyche." Kaj Sotala, "Advantages of Artificial Intelligences, Uploads, and Digital Minds," *International Journal of Machine Consciousness* 4.1, 2012, p. 1; Baum et al. argue in the same vein that "many experts do not consider it likely that the first human-level AGI systems will closely mimic human intelligence," Seth Baum, Ben Goertzel, and Ted Goertzel, "How long until human-level AI? Results from an expert assessment," *Technological Forecasting & Social Change*, 78, 2011, p. 191.
43. Hansen and Yudkowsky, 2013, p. 97.
44. Robin Hansen, 2013, "Whole Brain Emulation: Envisioning Economies and Societies of Emulated Minds," December 2, 2012, https://pswscience.org/meeting/whole-brain-emulation-envisioning-economies-and-societies-of-emulated-minds/.
45. Seth Baum "A Survey of Artificial General Intelligence Projects for Ethics, Risk, and Policy," Global Catastrophic Risk Institute Working Paper 17-1, 2017, p. 8.
46. "There are no metrics for measuring progress toward such ambitious forecasting targets [as high-level machine intelligence]." Ross Gruetzemacher et al., "Forecasting AI Progress: A Research Agenda," arXiv:2008.01848, August 4, 2020, p. 4; Ben Goertzel and Joel Pitt, "Nine Ways to Bias Open-Source AGI toward Friendliness," *Journal of Evolution and Technology*, 22.1, February, 2012.
47. Hannas et al., "China AI-Brain Research," (Center for Security and Emerging Technology, September 2020), https://cset.georgetown.edu/publication/china-ai-brain-research/, p. 42; Hannas et al, "China AI-Brain Research," p. 21.
48. Some institutes work in two or all three categories. In these cases we chose the research area most descriptive.
49. This study excludes multinational companies doing AI research in China, e.g., IBM, Intel, Microsoft, etc. A complete inventory would capture this research, which almost certainly expands China's own capabilities.
50. https://baike.baidu.com/item/百度研究院/13889005. Baidu also operates a Silicon Valley Artificial Intelligence Laboratory (硅谷人工智能实验室).
51. http://www.dlnel.org/lab/intro.html.
52. 打造通用人工智能的研究平台, https://www.tsinghua.edu.cn/info/1175/19900.htm. Baidu and Tsinghua University manage the NEL with support from other Beijing areas institutions. In August 2019, Baidu and the NEL stood up a Whampoa Academy (黄埔学院) aimed at cultivating industry's "chief AI architects." The group's name derives from Republican China's (1912–49) legendary academy that trained China's first modern military leaders, http://www.dlnel.org/news/details?id=185.
53. http://www.dlnel.org/study/study?cat_id=64.
54. Also called "PCL-Baidu-Wenxin" (鹏城-百度-文心).
55. A "parameter" is a configuration variable that adjusts during training. Its "connection weight" is analogous to the synaptic strengths between biological neurons, which also vary with learning. The human brain has some 86 billion neurons, 30 percent in the cortex. Neurons have on average some 7,000 synapses, yielding an estimated number of 100 to 500 trillion synapses in an adult human, https://en.wikipedia.org/wiki/Neuron.
56. https://www.marktechpost.com/2021/12/29/baidu-and-pcl-team-introduce-ernie-3-0-titan-a-pre-training-language-model-with-260-billion-parameters/.
57. The full name is Academy for Discovery, Adventure, Momentum, and Outlook, a forced expansion of "Damo." *Damo* is the spelling of 达摩, a Chinese word coined to transliterate Sanskrit *dharm* > dharma, hence the institute's alternate name "Dharma Academy;" Damo.alibaba.com/about.

58. In August 2017, some former Alibaba executives and technical experts formed a "general artificial intelligence company" (通用人工智能公司) called Qiyuan Shijie (启元世界, Qiyuan World) or Inspir.ai, with participation from Netflix, Tencent, NetEase, Tsinghua University, and the Hong Kong University of S&T, aimed at creating "intelligent agents with self-learning, self-decision, and self-evolution (自我进化) capabilities." The Beijing-based company received RMB 300 million (USD 47 million) in Series A funding in 2021. The name *Qiyuan* (literally: "start of an era") is also associated with "metaverse" (元宇宙), https://www.forbeschina.com/entrepreneur/56249; https://scc.pku.edu.cn/employment_22e903457bb0c5e8017be42a837e30cc_0.html.
59. https://www.infoq.cn/article/xix9lekuulcxewc5iphf.
60. https://www.bjnews.com.cn/detail/163635929714970.html.
61. https://www.bjnews.com.cn/detail/163635929714970.html.
62. Introductory remarks by Zhang Jianfeng, PKU and Princeton's E Weinan (鄂维南), and Wang Yuanzhuo (王元卓) of CAS's Institute of Computing Technology (中国科学院计算技术研究所), https://developer.aliyun.com/topic/download?id=8247.
63. https://www.tencent.com/en-us/about.html.
64. https://ai.tencent.com/ailab/zh/index.
65. https://baike.baidu.com/item/张潼/17743968. Zhang Tong attended Cornell and Stanford, and held high posts in IBM, Yahoo, and Baidu.
66. https://baike.baidu.com/item/俞栋/20727821. Yu was previously at MRA and is now directing Tencent's Seattle lab, which researches "the creative and decision making capabilities of AI." https://techseen.com/tencent-ai-lab-seattle-yu-dong/.
67. Literally "absolute comprehension;" https://syncedreview.com/2019/12/30/tencent-ai-juewu-beats-top-moba-gamers/.
68. https://ai.tencent.com/ailab/zh/news/detial?id=85.
69. https://www.pingwest.com/a/245807.
70. https://cloud.tencent.com/developer/news/319275.
71. https://ai.tencent.com/ailab/zh/news/detial?id=85.
72. https://www.huawei.com/en/news/2018/11/huawei-ascend-310-world-internet-conference-2018-award; https://e.huawei.com/cz/special_topic/event/2020q3/2020_09_cz_huawei_ai_webinar; https://www.cio.com/article/193320/mindspore-huaweis-all-scenario-ai-computing-framework-now-open-source.html.
73. https://www.163.com/dy/article/GB1J12NE0514K5HM.html.
74. Zengkun Li, "A Heuristic Search Algorithm Using the Stability of Learning Algorithms in Certain Scenarios as the Fitness Function: An Artificial General Intelligence Engineering Approach," Huawei 2012 Lab Central Software Institute Compiler and Programming Language Lab, December 8, 2017, https://arxiv.org/abs/1712.03043.
75. Zhang Xiaopeng (张晓鹏), "华为视觉计算预训练大模型，加速迈向通用人工智能" (Huawei Visual Computing Pre-trains Large Models to Accelerate towards Universal Artificial Intelligence), April 24, 2021, https://bbs.huaweicloud.com/live/HDC.cloud2021_live/202104241330.html.
76. Pan'gu is the mythical progenitor of the universe—the beginning and end of everything.
77. https://bbs.huaweicloud.com/forum/thread-123844-1-1.html.
78. PRC State Council. 国务院关于印发新一代人工智能发展规划的通知 (The New Generation AI Development Plan), SC 35, 2017.
79. "能够制造人工智能机器," literally: AI "able to make AI machines;" Donga Ilbo (Chinese edition), "华为：未来5年将投资1.5万亿美元构建人工智能生态系统" ("Huawei in the next five years will invest 1.5 trillion US dollars to build an artificial intelligence ecosystem"), January 15, 2020, https://www.dwnews.com/全球/60164854/韩媒华为和百度在争夺2030年人工智能的领头羊地位.
80. https://blog.51cto.com/u_14122493/3352658. The company, a competitor of Alibaba, is also called "Jingdong" (the full reading of its Chinese name). It was formerly called "360buy."
81. Donga Ilbo (Chinese edition). The description sounds like a recitation of the Turing Test.
82. https://arxiv.org/pdf/2110.13384.pdf.

83. https://www.sohu.com/a/472303775_114984. JD Explore Academy is headed by Tao Dacheng (陶大程), who also has an international pedigree, https://www.computer.org/profiles/dacheng-tao.
84. https://www.iyiou.com/interview/202107131019760.
85. https://www.pcl.ac.cn.
86. http://www.kejilie.com/leiphone/article/zI77n2.html.
87. http://static.scms.sztv.com.cn/ysz/zx/szdyxc/78287642.shtml.
88. https://www.sohu.com/a/505933408_115565.
89. CAE academician and Deputy Director of China's National Natural Science Foundation from 2013 to 2018, https://www.pcl.ac.cn/html/909.
90. https://www.sohu.com/a/463116712_115362.
91. Traditional AI "requires larger data and larger computing power as a basis to extract models, and then migrate and evolve. To learn from nature, how people think, and to solve these problems, we must find a more credible reference model in deep learning." "高文院士：强人工智能与类脑计算路线及安全对策" (Academician Gao Wen: Roads to strong AI and brain-inspired computing, and their safety measures), *Cyberspace Security Science and Technology* (网络空间安全科学与技术), May 26, 2021, https://www.secrss.com/articles/31459.
92. "Pengcheng Laboratory will carry out solid cooperation with the brain-inspired laboratory in artificial intelligence platform construction and brain-inspired application research." https://www.pcl.ac.cn/m/897/2019-07-03/content-2635.html.
93. https://baike.baidu.com/item/余凯/1383954.
94. Smaller Chinese companies claiming AGI engagement are Cyntek (辛特科技), founded in 2016 in Singapore with branches in Beijing and Shanghai. Cyntek is "committed to the theoretical exploration and application of AGI" and claims to have "an autonomous agent with general intelligence" (http://www.cynteksg.com/). Galaxy Eye Technology (北冥星眸) in Hangzhou stood up a GalaxyEye-AI Lab (北冥星眸通用AI实验室) in late 2020. Its goal is to build a "first-generation AGI prototype" (https://tech.ifeng.com/c/83JiRVV19Xl). IVFuture (智视科技) in Guangxi's Liuzhou High-tech Zone built a joint laboratory of "general artificial intelligence technology" in October 2020 (http://www.ivfuture.com/news_cont/page/304.html). The authors could not establish an expressed link between "general AI" and any of China's "Four AI Dragons" (AI四小龙)—Megvii (旷视), Cloudwalk (云从), Yitu (依图), and SenseTime (商汤), although the last may be on the path there. In January 2022, SenseTime announced a "super-large general AI infrastructure" (超大规模AI通用基础设施), https://www.geekpark.net/news/297626. The company has been researching large-parameter deep networks since 2017, https://zhuanlan.zhihu.com/p/392786166.
95. https://www.leiphone.com/category/industrynews/A9dgXfCrYmNn6jBl.html.
96. https://www.leiphone.com/category/industrynews/A9dgXfCrYmNn6jBl.html.
97. https://jobsearcher.com/j/research-scientist-general-ai-lab-at-horizon-robotics-in-cupertino-ca-gKlBal.
98. https://www.bigai.ai/index.html.
99. https://www.tsinghua.edu.cn/info/1182/83706.htm.
100. https://www.sohu.com/a/451788727_629135.
101. A Wuhan General Artificial Intelligence Platform was announced in May 2021 (see Subsection I).
102. A "Beijing General Artificial Intelligence Innovation Park" (北京通用人工智能创新园) is under construction in the city's Haidian District with completion scheduled for November 2024.
103. Hannas et al., "China AI-Brain Research"; Hannas and Chang, "China's 'New Generation' AI-Brain Project," PRISM 9.3, 2021; and Hannas and Chang, eds., *Chinese Power and Artificial Intelligence*. See also Andrew J. Lohn and Micah Musser, "AI and Compute: How Much Longer Can Computing Power Drive Artificial Intelligence Progress," (Center for Security and Emerging Technology, January 2022), https://cset.georgetown.edu/publication/ai-and-compute/.
104. Zhu's views on the shift to small data are a matter of record (http://www.ai.pku.edu.cn/info/1086/1913.htm). Deputy Director Dong Le (董乐), equally credentialed (https://baike.baidu.com/item/董乐/5884204), holds similar opinions, which she characterizes as AI's shift from "perceptual to cognitive intelligence." https://www.sohu.com/a/512316002_485557.
105. https://www.sohu.com/a/512316002_485557.
106. https://www.163.com/dy/article/GFBAPJ9V051481US.html.

107. 多模态人工智能产业联盟, http://news.cjn.cn/sywh/202112/t3884159.htm.
108. http://www.brainnetome.org/; See for example Zhang Tielin (张铁林) et al., "Self-backpropagation of synaptic modifications elevates the efficiency of spiking and artificial neural networks," *Science Advances* 7.43, October 2021.
109. "Pei" is his given name, "Wang" the family name, so "Pei Wang" follows western word order. East Asian languages, Chinese among them, invert the order, hence: 王培 (Wang Pei) in Chinese.
110. See "Selected Papers of Pei Wang," https://cis.temple.edu/~pwang/papers.html.
111. Pei Wang, *Non-axiomatic Logic: A Model of Intelligent Reasoning*, World Scientific, 2013; "Non-Axiomatic Reasoning System." NARS is featured in Seth Baum's "A Survey of Artificial General Intelligence Projects for Ethics, Risk, and Policy" (Global Catastrophic Risk Institute Working Paper 17-1, 2017). See https://home.x-in-y.com/pluginfile.php?file=/1069/course/section/591/纳思2016.pdf for a (Chinese) description. Wang's choice of near-homonym "纳思" (pinyin: nà sī, "hold or accommodate thought") to transliterate "NARS" renders the acronym; it is not a lexical item in Chinese. By happenstance, it coincides with the Chinese name of one of this paper's authors (韩纳思). The authors discovered this bizarre coincidence after reading Wang's work in English and, to stretch the metaphor, *accommodating* his *thoughts* to our framework. Also, both persons are Temple alumni.
112. See for example the June 19–25, 2017, course syllabus "通用人工智能的理论与实践" (Theory and Practice of General Artificial Intelligence) subtitled "NARS 系统解析" (NARS System Analysis). Lectures were mostly by Pei Wang and Liu Kai, https://home.x-in-y.com/course/view.php?id=34&lang=zh_tw. Liu Kai additionally manages a Chinese Artificial General Intelligence Society (中国通用人工智能协会) with Wang Pei and Xu Yingjin (see Subsection J) that sponsors the annual China AGI meetings, http://www.agi-society.cn/index.html.
113. http://www.hfitu.cn/newscontent/29835/.
114. https://www3.bhu.edu.cn/page/depart/jyxy/edu/js_info.asp?id=66.
115. BAAI was initiated as part of the Beijing Zhiyuan Action Plan (北京智源行动计划) also in 2018. Construction began under the same plan in November 2021 of a Beijing General Artificial Intelligence Innovation Park (北京通用人工智能创新园) in Beijing's Haidian district to be completed in 2024. The park will support a "general AI model computational platform and the industrialization of related technological achievements." http://zyk.bjhd.gov.cn/ywdt/rdgz/202111/t20211129_4497925.shtml?type=computer.
116. https://www.aminer.cn/research_report/5f44cbf13c99ce0ab7bc8db9?download=false.
117. "变革性、颠覆性突破" http://www.most.gov.cn/dfkj/bj/tpxw/201811/t20181123_143106.html.
118. Xu Bo is a BAAI director, director of CAS's Institute of Automation (CASIA), vice director of CEBSIT, and chair of China's "Next Generation Artificial Intelligence Strategic Advisory Committee."
119. The name "shenji" is composed of two morphemes meaning "neural" and "machine;" "In order to explore the possibility of strong artificial intelligence, the Zhiyuan Research Institute will successively develop Shenji 1 (a high-precision fine neural network simulation system), Shenji 2 (large-scale spiking neural networks), and Shenji 3 (a high-precision large-scale general-purpose intelligent simulation system);" The platform includes (1) CogNet "a database that integrates most types of human cognitive tasks," (2) BrainDB "a biological brain database" covering zebrafish to humans, (3) BrainPy "the first self-developed open source programming tool for computational neuroscience and brain-inspired computing in China," (4) DNNBrain "the first domestic toolkit for cross-study of deep neural networks and brain images," and (5) a "brain-like visual information processing model and algorithm library." https://mp.weixin.qq.com/s/rgau4_QZ-7J-iTa9c4aPXQ.https://www.baai.ac.cn/news_article?content_id=48&type=news.
120. 悟道要做的是让机器像人一样思考，迈向通用的人工智能, http://stdaily.com/index/kejixinwen/2021-06/01/content_1149556.shtml. The term *Wudao* translates to "road to awareness." BAAI (Zhiyuan) is collaborating with Tsinghua University in several projects to improve the training efficiency of super-parameter scale AI models and the "application of pre-trained models in biological research and Internet scenarios." https://posts.careerengine.us/p/61cd2e6c3b1d0824f1fe7695?from=latest-posts-panel&type=previewimage.
121. https://brain.bnu.edu.cn/.
122. Since 2011 and 2012, https://www.mcgoverninstitute.org/cn. CAS's Shenzhen Institute of Advanced Technology (中科院深圳先进技术研究院) teamed up with IDG-McGovern in 2014.
123. https://www.crossmodal-learning.org/home.html. Professor Zhang was elected to the German National Academy of Engineering in October 2021, https://en.acatech.de/person/jianwei-zhang-universitat-hamburg/.

124. https://www.ofweek.com/ai/2019-08/ART-201717-8120-30402700.html.
125. Xi-Nian Zuo et al., *Nature: Scientific Data*, December 2014. The paper, with 85 international co-authors, was funded by China's NNSF, National Basic Research Program (973), and National High Technology Program (863).
126. http://47.52.40.221/forum/node/2722.
127. www.cibr.ac.cn/.
128. https://jobs.sciencecareers.org/employer/513965/the-chinese-institute-for-brain-research-beijing-cibr-/.
129. https://baike.baidu.com/item/北京脑科学与类脑研究中心/22450631.
130. https://www.cibr.ac.cn/.
131. http://bcslab.ibp.cas.cn.
132. http://bcslab.ibp.cas.cn/yjgz/yjfx/.
133. http://www.cebs.ac.cn/en/duiwu_teacher.php?id=91; Tan served as CAS's vice director, deputy director of the Hong Kong liaison office, and is a senior CASIA leader, where he overseas biometric recognition and human surveillance. Pu, Xu, and Tan co-authored "Brain Science and Brain-inspired Intelligence Technology-an Overview" (脑科学与类脑研究概述), *Bulletin of Chinese Academy of Sciences*, 31.7, July 2016, an authoritative call to move beyond current AI, which "lacks generalizability" (缺乏通用性) toward BI systems "with autonomous learning capabilities."
134. http://bcslab.ibp.cas.cn/yjpt/gk/.
135. http://www.cebsit.cas.cn/yjsjj/zzjg/.
136. http://www.cebsit.cas.cn/yjsjj/zzjg/.
137. CEBSIT and CAS-ION share the same website (http://www.ion.ac.cn/) and seem to have merged—"CAS-ION" appears in parentheses after "CEBSIT" on the site's masthead, and both are under Pu Muming's directorship, CAS-ION since 1999. CAS-ION is also known as the Shanghai Institute of Neuroscience (上海神经科学研究所) from its original designation as "CAS, Shanghai Institute of Biological Sciences, Institute of Neuroscience" (中国科学院上海生命科学研究院神经科学研究所). Tracing these names and their lineages is one of Chinese databasing's most difficult and frustrating aspects (see below).
138. http://www.cebsit.cas.cn/yjsjj/szzc/.
139. http://stcsm.sh.gov.cn/xwzx/kjzl/20200706/56d8674bca674230b3b7699331d04d5f.html.
140. Follow-on funding for USD $2.85 billion runs through Dec. 2022, https://www.cingta.com/detail/17634.
141. Zhou Wenting, "Insights to Come from New Brain Research Base," *China Daily*, July 19, 2018. The complex includes the International Center for Primate Brain Research (国际灵长类脑研究中心) run by Pu Muming and Nikos Logothetis, who relocated his Max Plank Institute facility to China in 2020.
142. Hannas et al., "China AI-Brain Research," p. 30.
143. https://istbi.fudan.edu.cn/.
144. Feng Jianfeng chairs the Shanghai National Center for Mathematical Sciences and is dual-posted as a professor at the Centre for Scientific Computing and Computer Science at Warwick University, UK. He does single-neuron and neuronal network modeling, and causality analysis, https://istbi.fudan.edu.cn/lnen/info/1157/1629.htm.
145. www.ecst.com.cn/topic/detailIndex.html?id=1347; David Cyranoski, "Beijing Launches Pioneering Brain-science Centre," *Nature News*, April 5, 2018, https://www.nature.com/articles/d41586-018-04122-3.
146. http://www.pudong.gov.cn/shpd/department/20190228/019010005_9a5c450b-4981-4ed8-b62c-3321da01b438.htm.
147. http://bcmi.sjtu.edu.cn/index.html.
148. https://postd.sjtu.edu.cn/info/1004/1856.htm.
149. Lu Baoliang, Zhang Yaqian, and Zheng Weilong (郑伟龙), "情感脑机接口研究综述" (A survey of affective brain-computer interfaces), *Chinese Journal of Intelligent Science and Technology* (智能科学与技术), 3.1, 2021; Steffen Steinert and Orsolya Friedrich, "Wired Emotions: Ethical Issues of Affective Brain-Computer Interfaces," Science and Engineering Ethics, 26, 2020.
150. https://postd.sjtu.edu.cn/info/1004/1856.htm.
151. https://www.bsbii.cn/.
152. https://www.bsbii.cn/.

153. https://www.ljzfin.com/news/info/56264.html. Zhang is Vice President of CAS's Shanghai Branch and a CAS-ION researcher.
154. https://www.bsbii.cn/website/index.html. Lu Baoliang (BCMI), Guo Aike (CAS-ION), and Luo Qingming (Brainsmatics) are also on its staff.
155. Xu Zhang et al., "Brain science and technology: initiatives in the Shanghai and Yangtze River Delta region," *Nature Portfolio*, https://www.nature.com/articles/d42473-019-00213-5?proof=t.
156. https://www.bsbii.cn/.
157. http://www.leinao.ai/. Its partners include Fudan University; CAS's Shenyang Institute of Automation, Institute of Microelectronics, Institute of Electronics, and Institute of Neuroscience; Microsoft Research Asia, Baidu, iFlytek, and Datatang (数据堂).
158. In 2017, MOST created AITISA (新一代人工智能产业技术创新战略联盟) as a unified "artificial intelligence innovation ecosystem" under the New Generation AI Plan, i.e., to support government decision making and guide the distribution of funds. It has some 200 members. AITISA is led by PKU professor and Pengcheng Laboratory director Gao Wen (高文), and its secretary is Huang Tiejun (黄铁军), http://www.aitisa.org.cn/.
159. https://braindata.bitahub.com/.
160. Hefei CAS Brain-inspired Intelligence Technology Co. (合肥中科类脑智能技术有限公司), http://www.leinao.ai/about.
161. Not to be confused with the "Pengcheng" (鹏城) of Pengcheng Labs (see C.6 above).
162. https://baike.baidu.com/item/叶朝辉/4313061.
163. Brainsmatics has labs in Wuhan and Suzhou. Luo and his spouse Gong Hui also collaborate with CEBSIT, which has branches in both cities. Vivien Marx, "Qingming Luo," *Nature Methods*, 18.223, 2021, https://doi.org/10.1038/s41592-021-01091-w.
164. Brainsmatics is working with the Allen Institute for Brain Science on mouse brain connectomics. Vivien Marx, opcit.
165. http://www.brainsmatics.org.
166. I.e., "chip in a brain" and "brain on a chip." The authors include brain-inspired ("neuromorphic") chips with BCI in this topical area because of their experimental use as components of the interface.
167. Professor Dai wears several hats, indicative of the cross-disciplinary research done at the institute. He is a CAE academician, president of the Chinese Association for Artificial Intelligence (中国人工智能学会), director of the Broadband Network and Digital Media Lab in Tsinghua's Dept. of Automation (https://lixx1996.github.io/), and dean of Tsinghua University's School of Information Science and Technology (信息科学与技术学院). Dai has identified three areas where AI is poised for takeoff: in-memory computing (存算一体架构), neuromorphic computing (类脑计算), and photoelectronic intelligent computing (光电智能计算). See "Deep Learning Has Encountered Bottlenecks, Comprehensive Neural Observations Inspire the Next Generation of AI Algorithms," remarks at the 2020 International Conference on Computer Communication and Artificial Intelligence, reposted on the Tencent AI Discussion Forum, September 4, 2020, https://cloud.tencent.com/developer/article/1692341. Under Dai's tutelage, THUIBCS studies the brain at macro-, meso-, and microscopic levels.
168. The lab develops "next-generation core intelligent technologies such as brain-computer interface, human-computer interaction, and neuromodulation (神经调控)." This last area includes so-called "affective computing," https://brain.tsinghua.edu.cn/kxyj/yjfx.htm.
169. Zhu is vice director of China's National Engineering Laboratory of Deep Learning Technology and Applications (深度学习技术及应用国家工程实验室), which engages in "general AI" research (C.1 above).
170. Gao Shangkai is also a member of American Institute for Medical and Biological Engineering.
171. http://www.thubci.com/.
172. http://www.thubci.com/.
173. Artificial Intelligence Industry Alliance (中国人工智能产业发展联盟), "脑机接口技术在医疗健康领域应用白皮书" (White paper on the application of brain-computer interface technology in the medical and health field), July 2021.
174. https://www.cbicr.tsinghua.edu.cn.
175. https://www.tsinghua.edu.cn/kyzten/info/1003/1219.htm.

176. Jing Pei et al., "Towards artificial general intelligence with hybrid Tianjic chip architecture," *Nature* 572, 2019.
177. https://www.tsinghua.edu.cn/kyzten/info/1003/1219.htm.
178. Also called "computationally universal," i.e., a programming language that recognizes other data manipulation rule sets and can simulate their computational aspects, https://en.wikipedia.org/wiki/Turing_completeness; Youhui Zhang et al., 2020.
179. http://www.lynxitech.com/index.php/pcproduct; https://zhidx.com/p/156860.html. SynSense (时识科技), another neuromorphic chipmaker, claims similarly to implement "general purpose, real-time neural networks based on spiking neurons." https://www.synsense-neuromorphic.com/technology.
180. Distinct from Neuracle (博睿康科技), a Tsinghua BCI spin-off launched in 2011, http://www.neuracle.cn/sy; http://www.neuramatrix.com.cn/.
181. https://www.from-the-interface.com/China-BCI-neurotech/.
182. https://www.36kr.com/p/1491487908933765.
183. http://www.neuramatrix.com.cn/.
184. "人机混合智能" (human-machine hybrid intelligence). PRC State Council, "新一代人工智能发展规划" (New Generation Artificial Intelligence Plan), July 8, 2017; See Huey-Meei Chang and William C. Hannas, "Future paths and challenges," in Hannas and Chang, eds., *Chinese Power and Artificial Intelligence: Perspectives and Challenges*.
185. https://www.36kr.com/p/1491487908933765.
186. "The founding team has been conducting related academic research at the University of Pennsylvania and the Max-Planck Institute in Germany since 2010," https://www.36kr.com/p/1491487908933765.
187. http://news.tju.edu.cn/info/1003/46680.htm.
188. http://tj.teacher.360eol.com/teacherBasic/preview?teacherId=3191.
189. https://freewechat.com/profile/MzA5MDQyMTE4NA==?q=%40天大神经工程. Ming is executive director of the Tianjin University Research Institute (天津大学科研院) and holds several other university posts. He presided over seven NNSF projects, some military related, and has expertise in neural engineering and wave theory, https://baike.baidu.com/item/明东/10461571.
190. Liu Qian (刘茜) et al., "脑语者：脑科学时代正在走来" (Brain Talker: The Era of Brain Science Is Coming), 光明日报 (*Guangming Daily*), January 18, 2020, http://news.sciencenet.cn/htmlnews/2020/1/434964.shtm.
191. Liang Yu, ed., "China Unveils Brain-computer Interface Chip," *Xinhua*, May 18, 2019, http://www.xinhuanet.com/english/2019-05/18/c_138069590.htm.
192. https://www.sohu.com/a/315191613_99992021.
193. https://www.globaltimes.cn/page/202104/1222236.shtml.
194. Literally: "institute of brain-like chips and intelligent systems on a chip;" https://ibicas.fudan.edu.cn/16816/list.htm.
195. https://baike.baidu.com/item/闵昊/10853337.
196. https://www.yicaiglobal.com/news/china-first-wireless-brain-computer-interface-chip-for-animals-goes-on-display-in-shanghai.
197. The lab was established in 2002 under China's 10th Five-year plan (2001–2005), https://sbg.ecnu.edu.cn/32401/list.htm; Under Lin Longnian (林龙年), https://finance.sina.com.cn/tech/2021-01-10/doc-iiznctkf1304909.shtml.
198. Liangjian Lyu (吕良剑) et al., "A Fully-Integrated 64-Channel Wireless Neural Interfacing SoC Achieving 110 dB AFE PSRR and Supporting 54 Mb/s Symbol Rate, Meter-Range Wireless Data Transmission," *IEEE Transactions on Circuits* 67.5, May 2020.
199. https://asic-skl.fudan.edu.cn/.
200. http://news.sciencenet.cn/htmlnews/2020/12/450072.shtm.
201. https://sg.news.yahoo.com/shanghai-hospital-clears-ethical-reviews-082902452.html.
202. Formed in 2012 by three Shanghai Jiao Tong affiliates, including Lu Baliang's former student Cai Haoyu (蔡浩宇).

203. https://www.superpixel.com/article/15625.
204. https://www.tmtpost.com/5376907.html.
205. http://www.bbmi.zju.edu.cn/.
206. Duan, a 100 Talents selectee, trained at Kyushu University, University of Hawaii, and UC San Francisco. Duan was principal investigator at Shanghai's ION from 2000 to 2009. He was elected to CAS in 2007 and The World Academy of Science in 2008, https://person.zju.edu.cn/en/0006042#0; https://en.wikipedia.org/wiki/Duan_Shumin.
207. https://www.zhejianglab.com/. The lab is run by local party committee chair Zhu Shiqiang (朱世强). The first character of the lab's name 之, normally read "zhi" but "zhe" here, is an ancient pronunciation.
208. http://www.bbmi.zju.edu.cn/2021/0316/c43507a2268034/page.htm.
209. http://www.xinhuanet.com/english/2020-09/04/c_139342858.htm.
210. http://www2.scut.edu.cn.
211. http://www.ia.cas.cn/xwzx/mtsm/202006/t20200604_5602294.html.
212. http://www.ia.cas.cn/xwzx/mtsm/202006/t20200604_5602294.html.
213. The company's English name, visible on plaques posted online, avoids a literal translation of the Chinese name "South China Brain Control" (HNNK), to which it is commonly referred. The "i" presumably is added for iFLYTEK.
214. https://www.sohu.com/a/297786625_260616.
215. http://www.ihnnk.com/.
216. https://www.cheninstitute.org. The website has full English and Chinese versions.
217. https://www.cheninstitute.org/focus.
218. https://www.cheninstitute.org/nextquestion.
219. Tao Hu is also affiliated with the CAS Shanghai Institute of Microsystems and Information Technology (上海微系统与信息技术研究所). In his estimate: "It is not an exaggeration to say that the brain-computer interface is the main battlefield of the next intersection of life sciences and IT, and represents a new and potentially disruptive (具有潜在破坏性) technology field." http://www.xinhuanet.com/techpro/2021-06/03/c_1127524272.htm; Also called "Brain Tiger Technology" after the founder's given name 虎 "Tiger;" https://www.ithome.com/0/600/393.htm.
220. https://www.eet-china.com/mp/a103651.html.
221. http://www.ia.cas.cn/xwzx/mtsm/202006/t20200604_5602294.html.
222. https://www.technologyreview.com/2017/05/06/151950/mind-reading-algorithms-reconstruct-what-youre-seeing-using-brain-scan-data/.
223. http://www.nlpr.ia.ac.cn/cn/column/36.html. Tan Tieniu (谭铁牛), the only CAS academician among the lab's 44 researchers, is likely its senior figure, http://www.nlpr.ia.ac.cn/cn/column/36.html; http://www.ia.cas.cn/sourcedb_ia_cas/cn/iaexpert/201407/t20140717_4160162.html.
224. https://www.sohu.com/a/416529122_756411.
225. "2021-2022年度智源人工智能前沿报告." http://www.ncsti.gov.cn/kjdt/kjrd/202112/P020211231640762390337.pdf; http://www.ncsti.gov.cn/kjdt/kjrd/202112/P020211231640762390337.pdf, p. 106.
226. Thanks to CSET's Helen Toner for pointing out this important psycholinguistic feature.
227. PRC State Council, 国务院关于印发新一代人工智能发展规划的通知 (New Generation AI Development Plan), SC 35, 2017.
228. Problems associated with an intelligence explosion are a major recurring theme of publications sponsored by the Machine Intelligence Research Institute, https://intelligence.org/all-publications/.
229. Dr. Jason Matheny, personal communication, September 2020.
230. "AI and the Surveillance State," in Hannas and Chang, eds., *Chinese Power and Artificial Intelligence: Perspectives and Challenges.*
231. https://www.hrw.org/about/people/maya-wang.
232. Yuste, Rafael, Jared Genser, and Stephanie Herrmann, "It's Time for Neuro-rights: New Human Rights for the Age of Neurotechnology," in *Horizons*, 18, 2021. Yuste and Dario Gil, dialog with the authors, January 2021.

233. Hannas and Chang, "China's STI Operations," (Center for Security and Emerging Technology, January 2021), https://cset.georgetown.edu/publication/chinas-sti-operations/;中国科学技术信息研究所 (ISTIC), 甲子辉煌—中国科学技术信息研究所成立60周年纪念 *(60 Years of Glory—The 60th Anniversary of the Founding of the Institute of Science and Technical Information of China*), Beijing, 2016.
234. Hannas and Chang, "China's STI Operations," p. 41-44; H.R. 4747, "Open Translation and Analysis Center Authorization Act," introduced by Rep Joaquin Castro, July 28, 2021.
235. "Report of the National Commission for the Review of the Research and Development Programs of the United States Intelligence Community (unclassified version)," 2013. See also Hannas and Chang, "China's STI Operations," p. 1-7.
236. "Even the tiniest reduction of existential risk has an expected value greater than that of the definitive provision of any 'ordinary' good." Bostrom, 2013.
237. The project draws on open concepts implemented by members of CSET's cadre at USG institutions aimed at identifying foreign S&T research, capabilities, and breakthroughs that could impact global safety and U.S. security.
238. Hannas et al., "China AI-Brain Research."
239. These lists are shareable with the usual caveats.
240. Hannas et al., "China AI-Brain Research."
241. A subset of terms was chosen from the keyword glossary (Appendix) judged to be applicable to the problem with high certainty.
242. Including bottom-level entities—labs, university departments, ministry offices—if named in relevant citations or claimed by a performer as the affiliated unit.
243. Ross Gruetzemacher et al., opcit., p. 4.
244. Data sourced from Dimensions, an inter-linked research information system provided by Digital Science (http://www.dimensions.ai). All China National Knowledge Infrastructure content is furnished for use in the United States by East View Information Services, Minneapolis, MN, USA.
245. One needs to be mindful here and elsewhere of OPSEC ("operational security") constraints and the potential for data removal.
246. Linguist Morris Swadesh (1909–1967) created lists of basic concepts to track cross-language correspondences in comparative linguistics. Variants were proposed but the range generally fell between 100 and 200 terms.
247. "The Artificial Brain Laboratory was led by Hugo de Garis at Xiamen University. The project appears to have ended upon de Garis's retirement around 2010. Xiamen University now has a Brain-like Intelligent Robotic Systems Group, but this is not necessarily related to the ABL." Seth Baum, "A Survey of Artificial General Intelligence Projects for Ethics, Risk, and Policy," Global Catastrophic Risk Institute, *Working Paper* 17.1, 2017, p. 79.
248. McKenna Fitzgerald, Aaron Boddy, and Seth Baum, "2020 Survey of Artificial General Intelligence Projects for Ethics, Risk, and Policy, Global Catastrophic Risk Institute, *Technical Report* 20.1, 2020, p. 120.
249. http://www.ucmsr.com/public/uploads/20210419/b5fd69b2c44ada0d8faf70e6870e1bd0.pdf.
250. http://www.ucmsr.com/public/uploads/20210419/b5fd69b2c44ada0d8faf70e6870e1bd0.pdf.
251. http://www.ucmsr.com/public/uploads/20210419/b5fd69b2c44ada0d8faf70e6870e1bd0.pdf.
252. S&T Daily (科技日报) interview with Ji Rongrong, April 4, 2020. "学一个忘一个？人工智能遭遇'灾难性遗忘'解决方案治标不治本" (Learn One and Forget One? Artificial Intelligence Encounters 'Catastrophic Forgetting', and the Solution Treats the Symptoms but Not the Root Cause), http://www.kejie.org.cn/article/detail?id=1404213.
253. "腾讯发布AI开放平台以实现'通用人工智能'," (Tencent Releases AI Open Platform to Realize 'General Artificial Intelligence'), https://tech.huanqiu.com/article/9CaKrnKcNfu.
254. https://blog.csdn.net/Datawhale/article/details/117915062.
255. https://m.sohu.com/n/474201345/. Goertzel co-authored some 34 publications with Xiamen University-affiliated persons between 1991 and 2021.
256. Hou Yaqi (侯雅琦) Hou Xu (侯旭), *Science*, 373, August 6, 2021.
257. Zhu Jing (朱菁), Zhan Hao (詹好), Wang Ruoding (王若丁), "浅议AGI的定义和判定标准" (A Brief Talk on the Definition and Standard for Judging AGI), 2020, https://www.zhihu.com/zvideo/1323402777867595776.

258. https://informatics.xmu.edu.cn/info/1018/6312.htm; https://orcid.org/0000-0003-2946-6974.
259. Team members Jiang Min (江敏), Shi Xiaodong (史晓东), Zhou Changle (周昌乐), Zhou Jianyang (周剑扬), Shi Minghui (施明辉) still hold posts at Xiamen University, https://informatics.xmu.edu.cn/szdw/js/3.htm.
260. Hugo de Garis et al., 2009.
261. China's Mianyang (绵阳) nuclear weapons complex in Sichuan is the textbook example but there are many others.
262. The rapid jump from 18th to 5th place by Sichuan University College of Computer Science, another provincial outpost typically neglected in analyses of China's AI progress, is also noteworthy.
263. https://www.letpub.com.cn/index.php?page=grant. Dates are when the grant was made, not their period of execution.
264. "Omics" are "novel, comprehensive approaches for analysis of complete genetic or molecular profiles of humans and other organisms." "Omics Data," RD Connect, https://rd-connect.eu/what-we-do/omics/.
265. Luo Qingming is now at Hainan University (海南大学).
266. https://www.sciping.com/21461.html.
267. Now at Liaoning's Bohai University (渤海大学). See C.10 above.
268. https://blog.csdn.net/cf2SudS8x8F0v/article/details/110675316.
269. https://home.x-in-y.com/course/view.php?id=35.
270. http://agi-conference.org/.
271. http://m.xinhuanet.com/2021-06/01/c_1127518705.htm.
272. https://www.163.com/dy/article/GB1J12NE0514K5HM.html.
273. https://min.news/en/tech/010e995cff2998bcd858566a942da080.html.
274. https://www.sohu.com/a/510563870_115362.
275. http://multiverseaccordingtoben.blogspot.com/2009/07/will-china-build-agi-first.html.
276. Katja Grace, John Salvatier, Allan Dafoe, Baobao Zhang, Owain Evans, "Viewpoint: When Will AI Exceed Human Performance? Evidence from AI Experts," *Journal of Artificial Intelligence Research* 26, July 2018.
277. Grace et al., "Viewpoint: When Will AI Exceed Human Performance? Evidence from AI Experts," p. 734.
278. Hannas et al., "China AI-Brain Research."
279. https://www.sohu.com/a/512316002_485557.
280. https://www.tsinghua.edu.cn/info/1182/89554.htm.
281. https://www.ccf.org.cn/YOCSEF/Branches/Shenzhen/News/2019-10-24/669602.shtml.
282. Ming Dong, "Decipherer of 'Brain Language Code'," https://m.chinanews.com/wap/detail/zw/gn/2019/12-25/9042386.shtml.
283. http://webcache.googleusercontent.com/search?q=cache:E9cHElaRkfAJ: https://idea.cas.cn/viewconf.action?docid%3D79933&hl=en&gl=us&strip=1&vwsrc=0.
284. *Tencent Net* (腾讯网), October 31, 2019, https://new.qq.com/omn/20191104/20191104A0BD9U00.html.
285. *Qiushi* (求是), April 2019, http://www.qstheory.cn/dukan/qs/2019-02/16/c_1124114625.htm.
286. *S&T Daily*, June 20, 2019, http://www.stdaily.com/qykj/qianyan/2019-06/20/content_773207.shtml.
287. Part of the University of Electronic Science and Technology of China (电子科技大学). Yao is also director of the Sichuan Institute of Brain Science and Brain-inspired Intelligence (四川省脑科学与类脑智能研究院).
288. Ge Zhang (張鬲) et al. "Computational exploration of dynamic mechanisms of steady state visual evoked potentials at the whole brain level," Neuroimage 237, August 2021. The "digital twin brain" is reported as China's equivalent to the European Brain Project's "The Virtual Brain," i.e., a connectomics (brain modeling) project with potential BCI applications, https://www.sohu.com/a/474923303_162758.
289. https://www.jiqizhixin.com/articles/2022-01-07-7.
290. Helen Toner, "AI safeguards: Views inside and outside China," in Hannas and Chang, eds., *Chinese Power and Artificial Intelligence: Perspectives and Challenges*; https://futureoflife.org/beneficial-agi-2019/; https://www.baai.ac.cn/news/beijing-ai-principles-en.html.
291. *Guangming Daily* (光明日报), January 24, 2019, http:// www.cas.cn/zjs/201901/t20190124_4678012.shtml.

292. "The New Generation AI Development Plan," opcit.
293. https://www.berggruen.org/; CITIC Press Group (中信出版集团), Beijing, 2020. A shortened and translated version is now available in English: https://www.amazon.com/dp/9811623082?tag=philp02-20&linkCode=osi&th=1&psc=1.
294. https://www.berggruen.org/people/bing-song/.
295. https://www.sohu.com/a/401212638_114988.
296. https://www.sohu.com/a/407475303_260616.
297. "There is no abstract 'human' independent of the environment." https://www.sohu.com/a/401212638_114988.
298. Hajime Nakamura, *Ways of Thinking of Eastern Peoples: India, China, Tibet, Japan*, Honolulu: University of Hawaii Press, 1964; Richard E. Nisbett, *The Geography of Thought: How Asians and Westerners Think Differently ... and Why*, New York: The Free Press, 2003.
299. Beijing University Press (北京大学出版社), 2021.
300. https://www.sohu.com/a/147870097_260616.
301. Hannas and Chang, eds., *Chinese Power and Artificial Intelligence: Perspectives and Challenges.*
302. Ngor Luong and Ryan Fedasiuk, in Hannas and Chang, eds., *Chinese Power and Artificial Intelligence: Perspectives and Challenges.*
303. PRC State Council, 国务院关于印发新一代人工智能发展规划的通知 (New Generation AI Development Plan), SC 35, 2017.
304. PRC State Council, 国务院关于印发新一代人工智能发展规划的通知 (New Generation AI Development Plan), SC 35, 2017. "Basic Theory. 1. Big Data Intelligence Theory, Data-driven general artificial intelligence mathematical models and theories, etc." (数据驱动的通用人工智能数学模型与理论等).
305. https://zhuanlan.zhihu.com/p/21714160.
306. "Full Text of Xi Jinping's Report at 19th CPC National Congress," China Daily, November 4, 2017.
307. Jonathan Ray, "China at the nexus of AI and quantum computing," in Hannas and Chang, eds., Chinese Power and Artificial Intelligence: Perspectives and Challenges.
308. http://www.ia.cas.cn/dqyd/xxyd/202012/t20201223_5837241.html.
309. Xi Jinping, Speech to the 9th Politburo Study Session (中央政治局第九次集体学习), October 31, 2018.
310. Xi Jinping, Speech to the 9th Politburo Study Session (中央政治局第九次集体学习), October 31, 2018.
311. Dahlia Peterson, "Designing Alternatives to China's Repressive Surveillance State," (Center for Security and Emerging Technology, October 2020), https://cset.georgetown.edu/publication/designing-alternatives-to-chinas-repressive-surveillance-state/; Dahlia Peterson, "How China Harnesses Data Fusion to Make Sense of Surveillance Data," Brookings Institution, September 2021, https://www.brookings.edu/techstream/how-china-harnesses-data-fusion-to-make-sense-of-surveillance-data/; Dahlia Peterson, "AI and the Surveillance State," in Hannas and Chang, eds., *Chinese Power and Artificial Intelligence: Perspectives and Challenges*; Dahlia Peterson and James Millward, "China's System of Oppression in Xinjiang: How It Developed and How to Curb It," (Brookings Institution and Center for Security and Emerging Technology, September 2020), https://cset.georgetown.edu/publication/chinas-system-of-oppression-in-xinjiang-how-it-developed-and-how-to-curb-it/; Maya Wang's extensive bibliography is available at https://www.hrw.org/about/people/maya-wang#.
312. "I&W for projection of authoritarian power and/or infringement of ethical norms via collaboration, knowledge transfer, system deployment, or other means by China or other developers of AI and brain-inspired systems."
313. Hannas and Chang, "Chinese AI and neuroscience," in Hannas and Chang, eds., *Chinese Power and Artificial Intelligence: Perspectives and Challenges.*
314. Hannas and Chang, "China's 'New Generation' AI-Brain Project," PRISM 9.3, 2022.
315. Hannas and Chang, "China's STI Operations"; Hannas and Chang, "Unwanted Foreign Transfers of U.S. Technology: Proposed Prevention Strategies," (Center for Security and Emerging Technology, September 10, 2021), https://cset.georgetown.edu/article/unwanted-foreign-transfers-of-u-s-technology-proposed-prevention-strategies/; "Measuring risk and monitoring development," in Hannas and Chang, eds., *Chinese Power and Artificial Intelligence: Perspectives and Challenges*; and Tarun Chhabra et al., "Open-Source Intelligence for S&T Analysis," (Center for Security and Emerging Technology, September 2020), https://cset.georgetown.edu/publication/open-source-intelligence-for-st-analysis/.

316. Hannas and Chang, "China's STI Operations."
317. "A version of a product with just enough features to be usable by early customers who can then provide feedback for future product development." https://en.wikipedia.org/wiki/Minimum_viable_product.
318. Task shared with one of the linguist-collectors when the collect-data input balance shifts.
319. The local data repository should be accessible to non-linguist analysts and users; this condition ensures it.
320. The authors recommend one analyst be cleared TS/SCI to interact with the USIC and closed congressional committees.
321. Government announcements, enactments, budgets; statements by scientists, research directors, corporate spokespersons; organization manifestos, merger announcements, company claims; third-party discourse.
322. China curries international favor by participating in projects and fora on AI safety to build trust and encourage monitors to look the other way, i.e., the AI equivalent of running a biosafety unit with international cooperation, while weaponizing related pathogens in a nearby military lab, to cite a hypothetical example.

CSET CENTER *for* SECURITY *and* EMERGING TECHNOLOGY

CSET.GEORGETOWN.EDU | CSET@GEORGETOWN.EDU