国际与公共事务学院
School of International and Public Affairs

CIDEG

AliResearch
阿里研究院

中国城市治理研究院
CHINA INSTITUTE FOR URBAN GOVERNANCE

Open-Source AI:
The Value and Future of Cooperation

开源人工智能：合作的价值与未来

Research Report
研 究 报 告

2025.7.25

SHANGHAI.CHINA 中国.上海

CONTENTS

目录

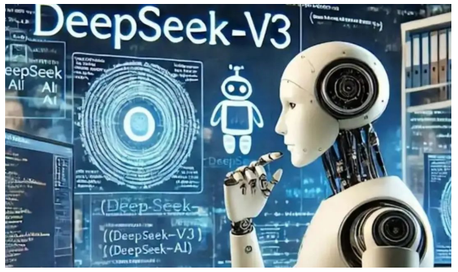

Chapter 1 Introduction: AI Enters the Open-Source Era

Artificial intelligence（AI）is entering an open-source era. Open-source AI projects are surging rapidly[①], and the performance of the most advanced open-source AI models is quickly closing the gap with their closed-source counterparts. At the same time, lightweight open-source AI models are dramatically reducing training and deployment costs, empowering every sector while meeting diverse, scenario-specific needs. Because of open source, more people are beginning to use, debug, learn from, and even build their own models. Because of open source, AI is going beyond the technology community to become a tool, teacher and partner for ordinary workers. Because of open source, our confidence that technological innovation can drive equitable progress for humanity is strengthened.

However, we cannot take for granted that open-source AI—or the positive outcomes it promises—will naturally materialize. In recent years, controversies surrounding the value and risks of open-source AI have intensified, to the extent that even the very definition of "open-source AI" has become a subject of significant dispute. Artificial intelligence—particularly large language models (LLMs) and large multimodal models (LMMs)—follows a technological and ecological logic distinct from that of traditional software, which means that many questions demand fresh responses and exploration: What is the true value of open source in advancing AI? How should we address the safety risks of open-source AI? How can we ensure its sustainable development? The governance reforms of open-source AI that are unfolding in the process of seeking these answers have thus become a "big issue," touching upon the future trajectory of AI's technological evolution and whether AI can genuinely contribute to the progress of human society.

Represented by China's Qwen, DeepSeek, MiniMax, Tencent Hunyuan and Baidu Ernie series of open-source large models, Chinese technical forces have become important contributors to the global open-source AI community. Open source has not only propelled China's domestic AI innovation and inclusive adoption, but also injected fresh impetus into the creation of global public goods and opened new avenues for international cooperation in AI development and governance.

① According to the latest "2025 Artificial Intelligence Index Report" by Stanford University's "Human-Centered Artificial Intelligence Institute", the number of open source artificial intelligence projects on the GitHub platform has increased from more than 1,500 in 2011 to 4.3 million in 2024, an increase of 40.3% in 2024.

This report offers a conceptual response to the "big issue" of the open-source AI era, grounded in China's practical experience with open-source AI development. We argue that open-source AI embodies the value of global cooperation beyond competition—a value that is not only essential for the technical production and reproduction of artificial intelligence, but also vital for enabling AI to address the real-world challenges of our time.

"Open source AI" refers to the production and application model of AI technology that publicly releases procedural elements such as AI-related software, tools, frameworks, models, data sets, etc. so that users can use or redevelop them. It is different from the closed source model that only provides AI products or services directly to users.

② For a discussion on global public goods, please refer to our team's research report "AI Security as a Global Public Good" released during the 2024 World Artificial Intelligence Conference: https://www.sipa.sjtu.edu.cn/Kindeditor/Upload/file/20241127/AI%20Governance%20as%20Global%20Public%20Commons.pdf

Chapter 2 The Value of Open-Source AI: Cooperation beyond Competition

AI was once regarded as a game for competitors, especially under the paradigm of the "Scaling Law", in which data, computing power, and network capacity became the core resources to be contested. Massive investments were poured into these areas, accumulating into the core competitive advantage of a handful of "players," who kept their gains tightly guarded. As the barriers to entry continued to rise, more and more participants were pushed out, leaving behind a small group of oligopolists struggling to deliver incremental breakthroughs amid increasingly fierce competition. Within this model, AI was framed as a "frontier technology" and evolved inside firms' laboratories into yet another instance of insular vanguardism—a closed, elite-driven pursuit detached from broader societal engagement.

But does competition truly reflect the evolutionary laws of AI? The release of DeepSeek V3 and R1 in late 2024 marked a milestone: more and more stakeholders realized that, while competition has value, it is not the only law governing AI's evolution. Real-world complexity challenges the effectiveness of the competitive model on at least three levels.

First, **the technological development paradigms of AI are multiplied**. The scaling law is only one paradigm but not the only one. Under resource constraints, high-performance models can still be achieved. Therefore, "competition" for resources does not automatically guarantee technological leadership.

Second, **the technology evolution models of AI are diversified**. AI's evolutionary bottlenecks are not merely about single-dimension resource input. They concern multi-dimensional, parallel linkages of knowledge (for example, the bottleneck for mathematical-reasoning models lies in formalized mathematical rule bases), of architecture (for example, the linear attention framework optimizes the original architecture) and of systems (for example, Geoffrey Hinton advocates extending learning from weight networks to hardware to integrate software and hardware). Competition thereby loses a clear object and anchor.

Third, **the application value of AI has to be realized in society**. AI's empowering value in advancing human progress cannot be realized in isolation within the firms' laboratory. Sustainable technical production and reproduction require AI to shift from a "frontier technology" to a "social technology." This transformation is reflected not only in the widespread integration of AI across diverse industries and sectors, but also in the embedding of social value objectives into the very trajectory of AI's technological evolution. As a result, the boundaries between domains become increasingly blurred, and the competitive model—predicated on clearly defined boundaries—loses both its foundation and its relevance.

It is based on this reflection that we argue for recognizing the positive role of open-source AI in overcoming the threefold bottlenecks of the competitive paradigm, and contend that its core value lies in opening up new pathways and new horizons for the future evolution of artificial intelligence through a novel model of "cooperation".

"Cooperation" as both a theoretical concept and a practice contains three dimensions: **relationality, consensus and recursion**. Firstly, relationality recognizes the interdependence among actors and emphasizes decisions based on interdependence features rather than individual attributes. Secondly, consensus acknowledges the existence of shared cognitions or interests among stakeholders while the value of cooperation lies precisely in uncovering and leveraging this common ground. Thirdly, recursion is that once relationality and consensus are appreciated, each stakeholder's decision logic shifts from a Nash equilibrium to a Kantian equilibrium—choosing whether to cooperate by applying the empathic principle similar to a Chinese proverb "do not impose on others what you yourself do not desire".

The threefold connotation of "cooperation" represents a breakthrough from the dilemma of the competition-based paradigm in AI's technological evolution. Recognizing relationality makes it possible to accept the potential for multi-dimensional and cross-boundary cooperation; recognizing consensus enables the activation of diverse driving forces to initiate such cooperation; and recognizing recursion allows these driving forces to be translated into concrete decisions to engage in cooperation. The underlying logic of open-source AI precisely reflects an acknowledgment and acceptance of these three dimensions—relationality, consensus, and recursiveness—and, on this basis, fosters cross-actor cooperation in the AI field. Such cooperation manifests across multiple dimensions and levels, including those between governments and communities, governments and enterprises, enterprises and enterprises, enterprises and communities, and among communities themselves.

It is important to note that although open-source AI, with cooperation as its core value, shares a connection with open-source software, it also exhibits fundamental differences that cannot be overlooked. Rooted in the idea of software as a form of freedom, the core value of open-source software lies in achieving the "four freedoms"—the freedom to use, study, modify, and redistribute the source code. Based on this foundation, cooperation in software open source has largely been an outcome, rather than a value-driven goal in itself. In contrast, for open-source AI, two critical distinctions arise. First, large-scale AI models still involve high thresholds for use, learning, modification, and redistribution, making it practically impossible for any single actor to fully exercise the "four freedoms." Second, innovation and application in AI do not easily generate positive feedback through these freedoms alone. Rather, cooperation becomes the key to overcoming developmental bottlenecks. In this context, cooperation evolves from a byproduct into a core value of open-source AI.

Chapter 3 Open-Source AI in China: Practices and Experiences

Open-source AI has, in theory, opened up new pathways and new horizons for the evolution of artificial intelligence. At the same time, it is gaining momentum in practice across countries worldwide. Within this global wave, China's open-source AI development—represented by models such as Qwen and DeepSeek—has attracted particular attention, injecting fresh ideas and new impetus into the global open-source AI landscape.

From a domestic perspective, China's open-source AI has already developed into a relatively complete and richly layered ecosystem. On the one hand, open source has become a basic consensus among AI stakeholders in China, with many actors choosing to open source their models, thereby forming a multi-dimensional, multi-sector, and multi-level open-source AI technology system. On the other hand, a relatively mature open-source community ecosystem has also emerged, covering key dimensions such as data, computing power, and models, providing effective support for a sustainable open-source AI landscape. Take ModelScope as an example: the platform has brought together over 16 million developers from more than 30 countries and offers over 70,000 high-quality AI models that are continuously updated. Leading Chinese models such as DeepSeek, Qwen, MiniMax, Wan were all first released on ModelScope, which has become a key window and channel for China's open-source AI to engage with the world.

From an international perspective, China's open-source AI community has become an important part of the global open-source AI community. The number of open-source AI projects from China ranks fourth only to the United States, Europe, and India.[③] At the same time, the performance quality of China's open-source AI projects has also attracted global attention: China's open-source AI projects have received more "likes" despite being fewer in number, and the average "likes" of each Chinese open-source AI project even ranks first in the world[④]. In terms of individual cases, the representative Qwen model has become the basic open-source model with the most derivative models in the world after October 2024. More generally, between 2019 and 2024, the number of open-source programmers from China ranked second or third in the world, only trailing to the United States or India.[⑤]

③ According to data from Stanford University's "2024 Artificial Intelligence Index Report".

④ According to the Stanford Artificial Intelligence Center AI Index report, the average "likes" is the total number of "likes" divided by the total number of projects.

⑤ Shengyu Zhao, et al.,"OpenRank Leaderboard: Motivating Open-Source Cooperations Through Social Network Evaluation in Alibaba,"Proceedings of the 46th International Conference on Software Engineering: Software Engineering in Practice, 2024.

开源人工智能:合作的价值与未来
Open-Source AI: The Value and Future of Cooperation

China's rapid development of its open-source AI ecosystem is both rooted in the global open-source movement and shaped by unique exploratory experience. This report summarizes that experience in four key aspects:

- First, **prioritizing accessibility by accurately identifying demand-side bottlenecks and laying a solid foundation for cooperative production**. Due to the complexity and opacity of AI models, open-source AI cannot easily replicate the "crowd-coding" development model of open-source software. In addition, the practical application of AI models requires a comprehensive support environment—ranging from compute resources and data to software-hardware integration. Through platform empowerment, the public's ability to use AI models has been significantly enhanced, enabling anyone to learn, apply, and interact with models in the most accessible ways. Ensuring inclusive access—allowing everyone to truly run the models—has become one of the defining features of China's open-source AI ecosystem development.

- Second, **upholding public value as a principle to incentivize participation from the supply side of open-source AI models**. At present, the high cost of AI development makes model contributors a critical bottleneck in the open-source ecosystem. Addressing their needs and building a full-chain service platform—from open-sourcing to deployment and use—is essential to encouraging broader participation. China's experience in this regard lies in embracing a public-oriented ethos: treating companies neutrally and equally, empowering small and medium-sized innovative enterprises, and building interconnected ecosystem nodes to promote the sustainable development of collaborative open-source AI.

- Third, **advancing inclusiveness as a direction to extend open-source AI into broader society and unlock collaborative value**. China's open-source AI ecosystem is not confined to technical communities. A growing number of model contributors have evolved from developers into creators and producers across diverse fields. Open source has created the possibility for ordinary workers to "play with" AI models, and the ecosystem has turned that possibility into a reality—taking concrete steps toward human-machine cooperation.

- Fourth, **embedding safety as a foundational baseline to promote the responsible development of the open-source AI ecosystem**. Amid ongoing debates over how open source may affect AI safety risks, China's open-source AI ecosystem has placed strong emphasis on reforming its safety governance frameworks. It has developed and implemented a series of innovative mechanisms for responsible governance, including AI ethics initiatives, baseline community governance principles, model content compliance reviews, behavioral codes of conduct for ecosystem participants, and third-party security evaluations of models. These mechanisms have effectively pooled the collective wisdom of developers, created a collaborative loop for risk control and quality improvement, and strengthened the ecosystem's overall safety resilience.

Chapter 4 Open-Source AI in China: The First Response from the Chinese Community

The practical experience of open-source AI in China shows the explorations and summaries of its pioneers, however, this does not imply that the controversies surrounding the value and risks of open-source artificial intelligence have been resolved. The differences between open-source AI and traditional open-source software prompt us to continue seeking answers to these contentious questions. Therefore, the research team, in cooperation with the ModelScope community, conducted a broad survey among Chinese open-source AI participants to capture the community's response to these debates. According to the team's observations, this is likely the first survey to reflect the views and perceptions of Chinese open-source AI practitioners. The survey was promoted on the ModelScope homepage and in developers' alliance groups from July 14 to July 19, and received a total of 1,136 valid responses.

Why did you join the open-source AI community? ——Consistent Recognition of open-source belief and philosophy

Among the five reasons, "agree with the open-source philosophy and believe that AI should be open-source" received the most support (70%). This result indicates that, despite various controversies, Chinese participants in the open-source AI community strongly support the open-source philosophy, viewing it as something that "ought to be" done rather than being solely based on practical or utilitarian considerations. Secondly, "facilitating easier use of models" garnered 43% support, reflecting the practical experience aimed at "accessibility", which indicates that the primary task of the artificial intelligence open-source community is to ensure the models are widely accessible. In contrast, reasons such as "Open-source AI model performance is leading", "receiving assistance in the open-source community", and "The open-source AI community offers a rich diversity of AI models" received only partial support, suggesting that the functional value of the open-source AI community is currently not the primary reason for the widespread participation of developers in open-source.

Why did you join the open-source AI community?

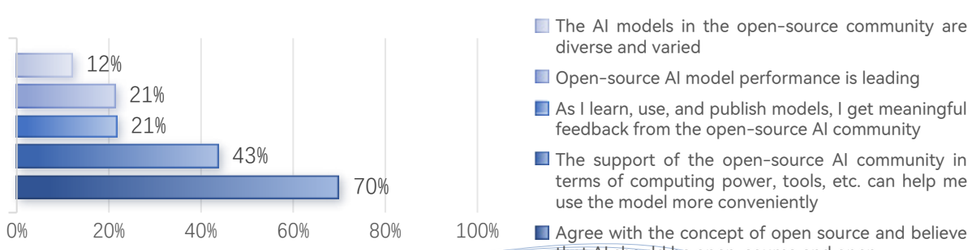

What is the value of open source for the development of AI? ——The value of open source is reflected in all dimensions

The questionnaire listed six potential benefits of open-source AI, all of which were endorsed by more than 60% of respondents (selecting "sometimes" or "greatly"). In contrast, the statement "open-source AI improves safety level" received the lowest support (61%), suggesting that the Chinese open-source AI community still harbors some reservations about whether open-sourcing affects AI-related safety risks. Moreover, when respondents were asked to indicate the highest level of support ("to a great extent"),"open-source AI improves safety level" (28%) and "open-source AI reduces costs" (34%) were chosen less frequently than"open-source AI boosts innovation efficiency"(42%), "open-source AI meets diverse needs" (41%), "open-source AI shortens time-to-market" (40%), and "open-source AI expands user scale" (38%). This pattern indicates that participants take a more pragmatic—or less confident—view of how substantially open-source can lower costs, acknowledging that the "last-mile" challenges of model engineering and real-world deployment remain significant.

What is the value of open source for the development of artificial intelligence?

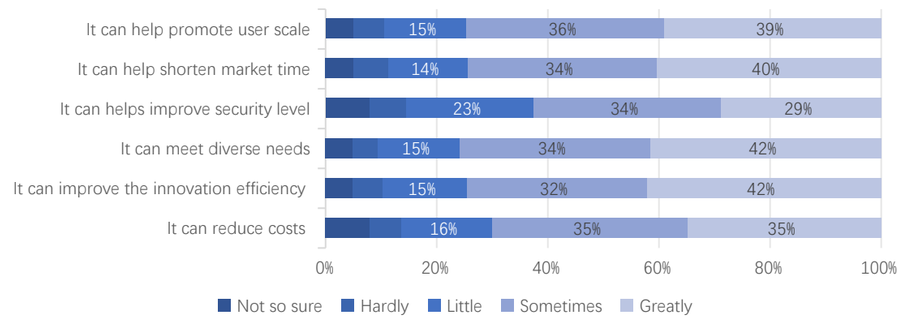

What is the value proposition of open-source versus closed-source AI? ——Convenience, cost, and risk considerations

A direct comparison of open-source and closed-source AI was a second focal point of the survey. Nine identical selection items were presented for both paradigms. When respondents chose open-source AI models, "with a low cost" (56%) and "easier deployment" (47%) were the dominant drivers, whereas "stronger safety control" (43%) was the leading reason for opting for closed-source solutions. This pattern positions open-source AI as the clear choice for cost-effectiveness and operational simplicity, while closed-source AI is perceived as the safer, risk-mitigated alternative. At the same time, "broader resource support", "better alignment with application needs", and "faster time-to-market" emerged as shared value propositions, each attracting more than 20% of selections. These findings demonstrate that both models can deliver competitive advantages in resource access, demand responsiveness, and commercial-

ization speed: high-impact open-source initiatives leverage vibrant communities for tooling and support, whereas established closed-source vendors provide specialized services and deep customization—each pathway enabling rapid iteration and swift market entry.

Why choose open or closed source AI?

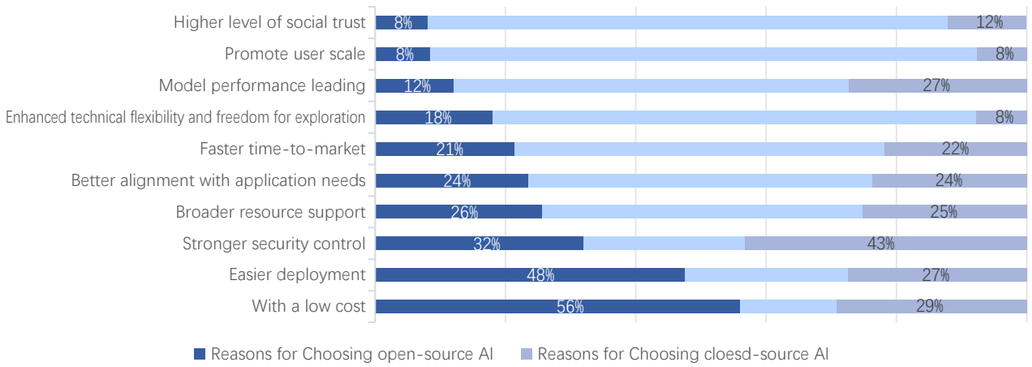

Is the value of open-source AI already a reality? ——Optimistic but with a long way to go

The survey asked participants how far their investments in open-source AI have translated into tangible benefits. Overall, Chinese open-source AI practitioners remain optimistic: nearly 80% of participants believe they will reap at least short-term gains. At present, 34% of participants report modest benefits, while 15% of participants describe the gains as substantial. Taken together, these figures confirm that the value of open-source AI is still unfolding—progress is evident, but the journey toward full realization remains long and demanding.

Is the value of open-source AI already a reality?

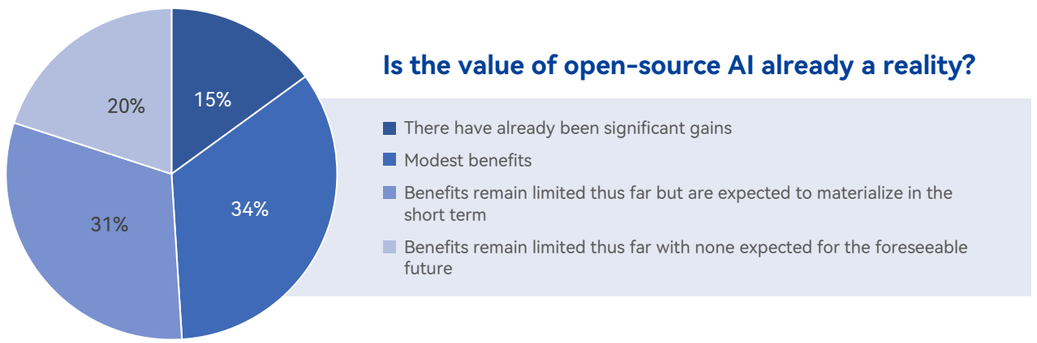

What are the factors that hinder the development of open source AI? ——The level of open source is not enough, and there are concerns about risks

The survey listed seven factors impeding open-source AI development. “Insufficient scope and depth of open source” drew the highest share of concern (36%), indicating that current openness levels still leave substantial room for expansion. In parallel, “unknown safety risks in open-source AI” (37%) outweighed “open-source AI has compliance risks”(23%), suggesting the community is more unsettled by safety uncertainties than by rule-based violations concern. Resource-related barriers—including “lack of long-term support” (25%), “shortage of open-source talent” (18%), and “inadequate compute resources”(16%) that all related to the resource constraints on open-source AI development, and in the future, we still need to increase support and investment in it.

What are the factors that hinder the development of open-source AI?

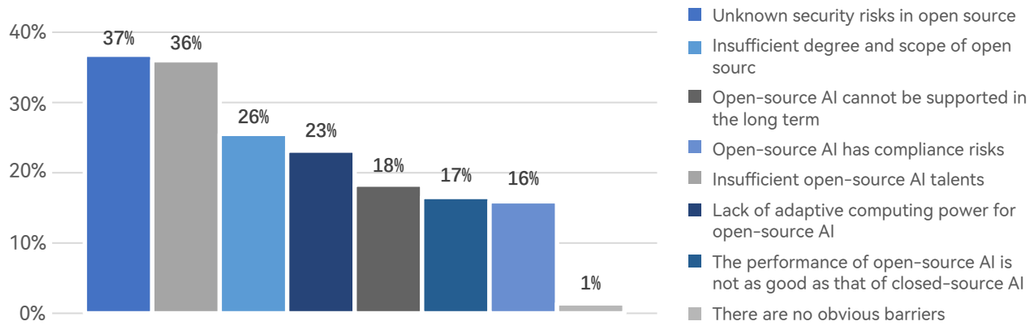

How will open-source AI evolve in the future? ——Embrace Open-Source AI But Care about the Safety Risks

The questionnaire presented six forecasts for the future evolution of open-source AI, each endorsed by more than 60 % of respondents (select “somewhat agree” or “strongly agree”). By contrast, the predictions that “open-source AI will become increasingly secure” and “open-source AI will evolve into an industry standard” attracted the highest shares of neutral responses (about 30% each), revealing that—amid widespread optimism—the Chinese open-source AI community still harbors reservations about advances in safety and standardization. Issues such as AI safety governance and industry standards intersect technology, ethics, society, and institutions in complex ways. Compared with the stronger consensus on more predictable trends like technological progress, application proliferation, or ecosystem cooperation, participants adopt a markedly more cautious stance when confronting these broader, longer-term objectives.

How will open-source AI evolve in the future?

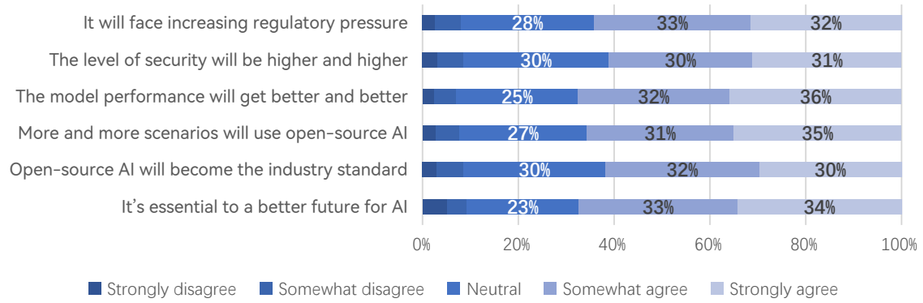

Chapter 5 An Action Agenda to Further Advance Responsible Open-Source AI

Open-source AI is still in its infancy. To better unleash its value and control its risks, this report proposes the following action agenda to promote sustainable, high-quality and inclusive development of China's and the world's open-source AI ecosystems.

I. Promote research and advocacy on open-source AI to sustain the open-source spirit and forge social consensus in the age of artificial intelligence. The industry-wide consensus established around open-source software will not automatically extend to the AI era. Debates over the concept, value, and risks of open-source AI persist—and are further complicated by geopolitical and other systemic factors. Recognizing the unique nature of open-source AI and clarifying its critical role in the future evolution of artificial intelligence require concerted efforts from all stakeholders. We call on diverse actors to jointly engage in research and public advocacy that illuminate the evolutionary logic of open-source AI ecosystems and help translate understanding into broad-based social consensus.

II. Champion reforms in open-source AI governance by aligning with the laws of technology and industrial development. Acknowledging the importance of open-source AI does not mean allowing its unchecked or disorderly growth. The fundamental differences between AI and software open-source—together with AI's entanglement with other complex governance challenges—make issues such as sustainability and risk controllability even more pressing. To put the open-source spirit into practice in the AI era, we call on all stakeholders to actively explore governance reform for open-source AI, pursue evidence-based consensus, and help establish an open, trust-based framework for cooperative governance.

III. Focus on the safety risks of open-source AI, and improve the safety level of the open-source AI ecosystem through system reform and organizational reform. Adhere to the concept of "AI safety as public good" and jointly build a cooperation-oriented safety governance ecosystem. Promote the socialization of open-source AI safety governance, and stick to the principle of innovating during applications, while ensuring safety during applications. Target at the challenges of changes in the safety governance environment of open-source cutting-edge AI models, and pay attention to bridge the safety capability gap. (Note: The first sentence should still be bolded)

IV. Broaden the frontiers of open-source AI by linking AI technologies with wider societal and industrial domains, and unleash the full potential of human-AI cooperation. While open-source software once inspired the creativity of technical communities, open-source AI

has the potential to stimulate the creativity of society as a whole. China's development experience shows that model applications and re-creations, supported by open-source services, are extending into all sectors of society and the economy, forming a broad and participatory collaborative ecosystem. To unlock this potential further, we advocate for stronger capacity-building efforts and for cultivating a new generation of interdisciplinary talent who combine AI expertise with deep domain knowledge—thus opening new avenues for human-AI cooperation.

V. Unleash the transformative significance of open-source AI by catalyzing multi-dimensional cooperation and fostering an ecosystem that is co-built, co-existed, and co-shared. Technology and institutional arrangements often co-evolve. Just as the "Scaling Law" once shaped a paradigm of competition, the rise of open-source AI requires the construction of a cooperation-oriented institutional environment. We therefore call on all stakeholders to collaboratively build such an environment—one that fosters cooperative values, ethical norms, operational mechanisms, and equitable benefit-sharing—to embrace the era of open-source AI and the broader future of collaborative artificial intelligence.

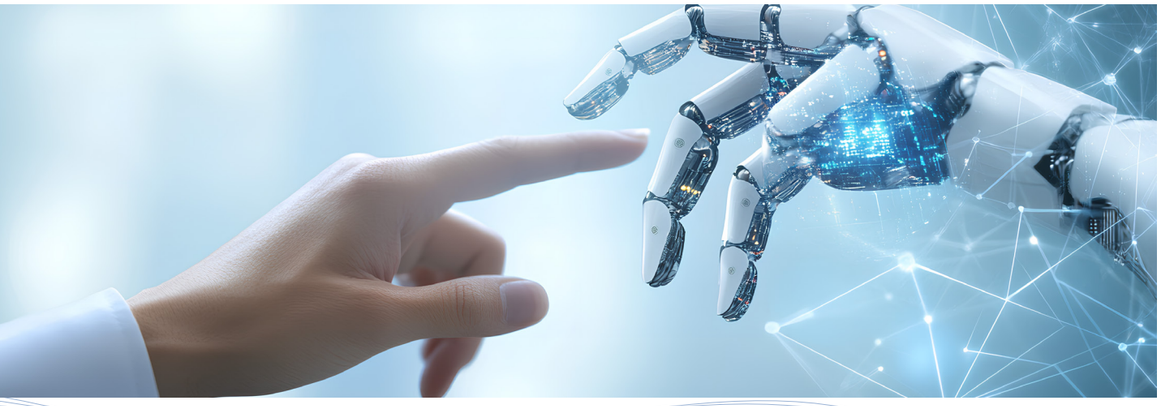

Guiding Organizations:

Shanghai Municipal Economy and Informatization Commission

Shanghai Cyberspace Administration

Supporting Organisations:

China IGF

OpenAtom Foundation

China–BRICS AI Development and Cooperation Center

ModelScope Community

China Association for Science of Science and Technology Policy Research, Frontier Science and Technology and Industrial Innovation Governance Committee (Preparatory Committee)

Mentoring Advisors:

Xue Lan	Distinguished Professor of Arts, Humanities and Social Sciences, Dean of Schwarzman Scholars, Tsinghua University
Cui Zhiyuan	Professor, School of Public Policy and Management, Tsinghua University
Wu Jiannan	Chair Professor, Shanghai Jiao University
Huang Qixuan	Professor, School of International and Public Affairs, Shanghai Jiao Tong University

Report Drafting Group:

Jia Kai	School of International and Public Affairs, Shanghai Jiao Tong University
Zhao Jing	School of Public Policy and Management, Tsinghua University
Fu Hongyu	AliResearch
Peng Jingzhi	AliResearch
Wang Zhe	Beihang University
Gao Longxu	School of International and Public Affairs, Shanghai Jiao Tong University
Tan Bohao	School of International and Public Affairs, Shanghai Jiao Tong University
Liu Boyan	School of Public Policy and Management, Tsinghua University