Simon Institute for Longterm Governance

China's AI Governance: A Conversation with Professor Zhang Linghan

March 2026

The Simon Institute (SI) is publishing periodic conversations with leading legal and policy experts to explore how different countries and regions are navigating the governance of AI, and how emerging domestic practices might inform international AI cooperation.

In the first of the series, SI spoke to Professor Zhang Linghan about China's approach to AI governance. Zhang is Director of the Institute of AI Law and Governance, China University of Political Science and Law, and was a member of the United Nations (UN) Secretary-General's High-Level Advisory Body on AI. She is an expert member of committees at China's Ministry of Industry and Information Technology and Ministry of Public Security and has participated in the drafting of many laws and regulations related to AI algorithms, data protection, and platform governance. In 2024, she led the drafting expert group of the China AI Law (Scholar Proposed Draft). The same year, she was named to TIME's list of the world's 100 most influential people in AI.

This interview was conducted in Mandarin Chinese and translated to English. It has been edited for length and clarity. The views expressed are those of the interviewee and do not reflect the official policy or position of SI.

Contents

Evolution of Chinese policymaking for AI 3
1. Trajectory since 2017
2. Priorities and points of emphasis

Chinese policy and regulatory processes 4
3. Drafting and consultation
4. Ensuring adaptability
5. AI+ initiative

Outlook for domestic AI governance 6
6. Future priorities for AI legislation
7. Frontier AI regulation
8. AI Safety Governance Framework 2.0
9. Anthropomorphic AI regulations

International context 8
10. Comparisons with EU regulatory approaches
11. World AI Cooperation Organization

Cover image courtesy of Professor Zhang Linghan.

Evolution of Chinese policymaking for AI

1. **How would you describe the trajectory of China's AI governance since 2017, when China's Next Generation AI Development Plan was released?**

I published a paper that details the overall development trajectory of China's AI governance since 2017. The legal governance of AI in China has evolved through three distinct stages, moving from foundational exploration to focused regulation, and now into a phase of system integration. The current focus is on achieving a synergy between high-quality development and high-level security through a hierarchical legal structure, at the apex of which is a dedicated AI Law.

Exploration stage (2017–2021). China established a regulatory system for AI inputs such as network infrastructure and data. Technology governance was mainly based on ethical norms and treated applications as the main object of governance.

Focused regulatory stage (2022–2023). The governance framework became more specialized, prioritizing information content security and establishing a tiered regulatory system specifically for Generative AI.

System integration stage (2024 onwards). Guided by the 20th Communist Party of China Central Committee's directives to "improve the mechanisms for developing and managing generative artificial intelligence" and "institute oversight systems to ensure the safety of artificial intelligence," China is accelerating the development of a multi-layered legal architecture. Central to this system is the formulation of dedicated AI legislation, supported by administrative regulations, international coordination, and technical standards to ensure the healthy, secure development of AI.

2. **What are the priorities and points of emphasis of China's approach to AI governance?**

Although specific points of emphasis may evolve over time, the core priorities of China's AI governance framework have remained relatively stable over the past seven to eight years.

A first priority is the protection of information integrity, a concern that is not unique to China but is shared across jurisdictions. Notably, when the UN Secretary-General's High-Level Advisory Body on Artificial Intelligence conducted an expert survey on AI-related risks, harm to information integrity emerged as the issue of greatest concern among respondents.

A second set of priorities is more closely linked to the institutional features of China's political and socialist system, particularly the protection of labor rights. One of China's earliest regulatory instruments in the AI domain, the Management Provisions on Algorithmic Recommendation Services, explicitly requires the protection of workers' rights when algorithmic systems are used for the scheduling and management of gig workers. The regulatory logic here is that inadequate protection of labor rights in algorithmic management contexts could pose risks to broader social stability.

A third priority, which aligns closely with regulatory approaches in other jurisdictions, is consumer rights protection. In China, recommendation algorithms were first widely deployed on e-commerce platforms, and many aspects of daily life, including retail consumption and food delivery, have become deeply integrated with online platforms. Against this background, regulators have placed particular emphasis on preventing unfair trading practices, including those arising from algorithmic pricing mechanisms. In this respect, China's regulatory concerns closely resemble those addressed by agencies such as the US Federal Trade Commission.

Beyond these relatively stable priorities, regulatory focus continues to adjust in response to technological developments. One notable example is the increasing attention paid to extreme risks associated with advanced AI systems. Beginning last year, relevant government departments and expert groups involved in drafting China's AI Safety Governance Framework 2.0 elevated such risks to a more prominent position within policy discussions, reflecting a broader shift in how advanced AI risks are conceptualized and addressed.

Chinese policy and regulatory processes

3. **Can you tell us about the process of drafting AI regulations and consulting with industry and academics?**

There isn't a single standard process—the legislative landscape in China is categorized into a hierarchy consisting of national laws, administrative regulations, departmental rules, etc., each governed by distinct procedures under the Legislation Law of the People's Republic of China.

National laws are enacted by the National People's Congress and the Standing Committee of the National People's Congress in accordance with legally prescribed authority and procedures.

The departmental rules are generally issued by Departments of the State Council or other statutory institutions (e.g., the Cyberspace Administration of China (CAC), the Ministry of Industry and Information Technology) to address specific sectoral issues. Most current AI regulations in China, such as those for algorithmic recommendations and deep synthesis, fall into this category. These rules are formulated and issued by the relevant departments, so the process is much faster—approximately six months to one year, compared to at least two to three years for a national law.

Opinions from industry and academia are collected at many stages. Departments tend to establish a drafting group with experts from academia and sometimes industry. Alongside consultation with other bureaus in the department and subsequent legal review, consultations are held with external experts (e.g., in law, journalism, and industry). Drafts are then released for public comment, reflecting the principle of "democratic lawmaking" (民主立法).

By way of example, after the public consultation for the Interim Measures for the Management of Generative AI Services had already begun, I joined forces with other experts to publish a proposal for five specific amendments and sent it to the official inbox for the consultation. Later, I was honored to see that in the final official version, four of the suggestions had been formally adopted.

4. **How do Chinese regulators strive to ensure that regulation remains appropriate given the pace of AI innovations?**

This is a difficult problem faced by legislators all over the world. My personal feeling is, if a law lacks adaptability, it's usually not because the harmful results were different from expected, but rather that the governance methods lack adaptability. I'll give a simple example: the kinds of harmful content we want to guard against—such as fake news or pornographic material—remain constant regardless of the technology used to produce it. The rights the law intends to protect are also relatively stable.

But if you have a specific regulatory route targeting a specific technical approach, it might not be sufficiently adaptable. I have a very vivid example. The 2021 Algorithmic Recommendation Regulation implied that user tags are used to conduct content recommendation, reflecting the practices of many service providers using older technology. Several technologically advanced platforms have noted that their algorithms utilize collaborative filtering rather than explicit user tags. While existing regulations cover most service providers, they often fail to account for the full spectrum of these diverse algorithmic methodologies.

To ensure adaptability, China has adopted a cautious approach toward high-level AI legislation, which explains the ongoing debate surrounding the national AI Law. Many Chinese scholars argue that the European Union (EU) AI Act's rigid framework has stifled innovation—a fate they wish to avoid. Consequently, China currently prioritizes departmental rules and policy directives, as these can be updated frequently to keep pace with rapid technological shifts. Another example of how Chinese regulators have demonstrated flexibility is by streamlining filing procedures for generative AI services, while maintaining the same fundamental expectation for transparency and traceability.

5. **The Guiding Opinions on Implementing the "AI+" Initiative (hereafter "Guiding Opinions") released in 2025 are one of the most significant developments for encouraging and shaping AI adoption in China. How are central and local government authorities trying to implement these?**

The AI+ initiative is a significant measure for AI application. On the one hand, China is striving to catch up with the world's frontier models; on the other hand, it is also leveraging its unique advantages in the field of AI. In China the use cases for AI application are very broad, the population is relatively large, and data is relatively abundant. These are some of China's strengths and advantages in developing AI applications.

The implementation of China's AI+ strategy follows a top-down model where central guiding opinions trigger a cascade of specific action plans from Departments of the State Council and local governments. This approach ensures that while the central government sets the strategic direction, individual departments and localities tailor implementation to their specific industrial layouts and local characteristics.

For example, the National Health Commission released implementation opinions focusing on "AI+ Healthcare," aiming to establish high-quality medical datasets and vertical LLMs for clinical specialties by 2027. The Ministry of Transport issued a plan to integrate AI across the construction, management, maintenance, and operation of transport infrastructure, targeting widespread application in typical scenarios by 2027.

All sub-plans must maintain strict alignment with the Central Government's "Guiding Opinions." Each city (Beijing, Shanghai, Shenzhen) offers different incentives based on its local characteristics, such as Beijing's focus on foundational scientific AI research. The integration of AI+ plans into the 15th Five-Year Plan ensures that these initiatives will remain national priorities through at least 2030.

Outlook for domestic AI governance

6. **A draft AI Law was listed in the State Council's Legislative Plans for 2023 and 2024, but in 2025, the reference was replaced with the phrase "advance legislative work for the healthy development of artificial intelligence." You have suggested that this reflects the need for a system-wide approach that not only includes an overarching AI law but also integrates legislation relating to key inputs and specific sectors. In your view, what should be the top priorities to advance first within this broader approach to AI legislation?**

We legal scholars generally think of four categories of legislative task: enactment of new laws, amendment of laws unsuited to current technology industry developments, abolition of unsuitable laws, and interpretation of existing legal articles (立改废释).

When existing legal frameworks prove insufficient to address the challenges posed by artificial intelligence, legislative priorities should be calibrated toward areas in which concrete social disputes have already emerged and normative uncertainty has become practically consequential. In such contexts, the function of legislation is not to speculate abstractly about future risks, but to provide definitive legal responses to real-world cases that cannot be adequately resolved under current law. A frequently cited illustration is traffic accidents involving autonomous vehicles: within traditional liability regimes premised on human fault or control, the attribution of responsibility may be indeterminate or structurally incoherent. As disputes of this kind increasingly arise in practice, the absence of clear legal rules undermines legal certainty and dispute resolution. Accordingly, developing an appropriate legal framework to address such scenarios constitutes a pressing legislative task, driven less by technological novelty per se than by the immediate need to maintain the operability and credibility of the legal system.

At the same time, there is a broad consensus that high-level legislation for AI will inevitably be introduced. Relevant departments announced at our forum at the World AI Conference last year that they are advancing this. Why is high-level legislation indispensable? It's because currently, China's AI governance mainly relies on departmental rules and policies, and the authority of these is lower than existing laws. This means they cannot conflict with existing laws. There must be a new high-level AI law to resolve these conflicts. For example, the tension between copyright protection and the use of copyrighted material in training datasets necessitates resolution through comprehensive legislation at the national or international level.

7. **How are Chinese regulators thinking about whether dedicated regulatory approaches are needed for frontier AI?**

This issue is currently being jointly studied and discussed by relevant stakeholders. First, one characteristic of frontier AI is that there are not many institutions in the world that truly possess the relevant capabilities. Therefore, this is a highly consequential issue that has, to date, attracted relatively limited attention.

Second, frontier AI involves technology that is so sophisticated that it cannot be addressed by universal technical standards. At this stage, governance mainly relies on the ethical conduct of frontier AI researchers and timely reporting to the government.

Third, at the legal level, it is difficult to establish effective regulatory mechanisms. I think currently, government departments probably don't have very efficient methods, apart from requiring regular red-teaming tests and regular information sharing. We have to admit that few experts know how effective these methods are. Technical capability in government is also a challenge, and not just in China. For example, in 2023, the Biden administration proposed that for frontier models, the results of red-teaming tests should be shared with the government in a timely manner. The question is: even if companies shared the results with the government, would the government have sufficient technical personnel to understand them?

In China, regular red-teaming is not written into law but is part of the "AI Safety Governance Framework 2.0." This indicates that the Chinese government attaches great importance to the extreme risks posed by artificial intelligence and has already made corresponding arrangements in relevant policy documents. China's scientific community is also paying close attention to these issues and has undertaken a number of important testing and experimental efforts.

8. **The AI Safety Governance Framework 2.0 (hereafter "the Framework") gives a detailed overview of AI risks - including content security, data security, loss of control and facilitating weapons development - and recommends measures for mitigating them. Can you tell us more about the nature of this document and how you expect it to be used?**

You can understand it as a roadmap or guidelines for AI safety governance. It is not a mandatory regulation or a policy, but a blueprint for the entire AI safety ecosystem. Towards the end of the document there are specific behavioral guidelines for developers and deployers.

Do you think there might be more specific standards or regulations in the future to implement some of the things inside it?

This Framework is already the most granular and actionable document to come out on AI governance. It explicitly tells people how China perceives AI safety. Rather than talking about its many measures, I think the core information that can be read from it is how Chinese regulators perceive AI safety. I believe future developments, whether laws and regulations or technical standards, will unfold based on the perceptions reflected in the Framework.

Can you tell us more about how the Framework was drafted?

The Framework was written by the Cybersecurity Association of China under the guidance of the CAC. The association coordinated relevant departments and teams in China, including China's most cutting-edge AI safety technical teams and policy experts, to draft the document *(Editor's note: participating organizations are listed on page 82 of the Framework).*

Is the audience domestic or international?

Both, but if I had to pick one, domestic. The most important part was everyone coming together and aligning on the idea that AI safety is not just content safety or data security but has many constituent elements.

9. **CAC recently issued draft regulations for anthropomorphic interactive AI. What are the main provisions and how do they compare to similar regulations in other jurisdictions?**

The draft regulation requires service providers to have the capability to identify and respond to users' extreme emotions and to set up special protection measures for minors and the elderly. These are mostly standard requirements for emotional companionship and anthropomorphic interaction services. The US also has similar requirements, for example, with extreme emotions like suicidal intention, there has to be human intervention, you have to refer the user to crisis service providers. *(Editor's note: The references here to the US refer to state legislation in California and New York.)* To respond to dependency risks, China requires a prompt every two hours [to remind users to take a break], the US requires a prompt every three hours. I think there is no essential difference here.

There are other important differences though. The draft regulation requires special protections for both minors and the elderly, but US regulations only target minors. Following the public consultation process, it is to be expected that there will be further revisions in many respects.

International context

10. **Apart from China, one of the most active government regulators of AI is the EU. What are the main similarities and differences between the regulatory approaches in the EU and China?**

I mentioned earlier (see Question 2) that there are points of divergence. China does place importance on rights protection, but its conception of rights is more collectivist in orientation. In general, regulatory priorities tend to emphasize the protection of the rights and interests of working people at the lower end of the social and economic structure, as well as consumer rights. By contrast, the EU places comparatively greater emphasis on individual fundamental rights as a core normative reference point.

With respect to regulatory approaches, the EU AI Act appears, at least at the textual level, relatively stringent, particularly in its extensive standards for general-purpose AI and so-called high-risk AI systems. However, a closer examination suggests that the actual number of enterprises that fall within the scope of enhanced transparency obligations or registration requirements for high-risk AI may be relatively limited. In contrast, China's AI regulatory framework applies more uniformly across AI service providers. In practice, Chinese AI companies are generally subject to comparatively high compliance standards and are required to bear substantial compliance costs regardless of scale or risk categorization.

At the same time, there are notable areas of convergence. One example is the regulation of AI-generated content through labeling requirements. The EU has proposed labeling obligations, and jurisdictions such as California impose labeling requirements once platforms exceed a certain user threshold. China, however, adopts a more comprehensive approach: labeling obligations apply broadly across contexts, and corresponding technical standards and testing standards have already been established.

From a longer-term perspective, this suggests that the underlying regulatory logic of China's AI governance framework exhibits a higher degree of stability and regulatory autonomy. Moreover, attention must be paid to the distinction between formal legal texts and actual implementation. In China, once regulatory rules are formally adopted, they are generally translated into concrete compliance obligations for enterprises. By contrast, although the EU AI Act has undergone a lengthy legislative process—marked by repeated revisions, debates over simplification, and delays—the extent to which its provisions have been fully implemented in practice remains uncertain at this stage.

11. **In July 2025, Premier Li Qiang proposed the launch of the World AI Cooperation Organization (WAICO). What will it focus on and how do you envisage it complementing the role of other existing multilateral organizations?**

WAICO is primarily oriented toward promoting international cooperation in artificial intelligence. Capacity-building has long been a central pillar of China's approach to international AI cooperation, and WAICO functions as an important complement to existing mechanisms within the UN framework by helping to advance capacity-building initiatives.

Within the broader framework of AI capacity-building cooperation, China seeks to contribute foundational infrastructure, including data resources, data centers, and AI models. At the same time, China also aims to draw on its comparative strengths in AI applications by offering diverse and effective use cases in cooperation with partner countries.


Simon Institute for
Longterm Governance