July 15, 2026

EXPERTS TAKE

"Why we should cooperate with China on AI regulation"

An Interview with Tom Abram

Tom Abram is a Research Fellow at the Foundation for Strategic Research (FRS). His work focuses on the foreign and security policy of China and Taiwan, strategic stability and deterrence in Northeast Asia, and the intersection of emerging technologies and national security. Abram is also a PhD candidate at the Centre for International Studies (CERI) at Sciences Po/CNRS under the supervision of Dr Stéphanie Balme, where he examines China's role in the emerging global governance of artificial intelligence. He holds an MSc in China in Comparative Perspective from the London School of Economics (LSE) and an MPhil in International Relations from Sciences Po-Paris.

Foreword

Maud Descamps

Artificial intelligence has rapidly emerged as one of the defining geopolitical, economic, and technological issues of the twenty-first century. Beyond transforming economies and societies, AI is increasingly shaping global power relations, industrial competitiveness, information control, and national security. As governments race to develop and regulate these technologies, the question of who defines the rules, standards, and norms governing AI has become a matter of growing strategic importance. In this context, China has positioned itself not only as a leading AI power, but also as an increasingly influential actor in global debates surrounding AI regulation, ethics, and technical standards.

China's approach to AI governance must be understood within the broader context of its strategic planning framework and long-term technological ambitions. Artificial intelligence was identified as a key priority in the 14th Five-Year Plan (2021–2025), which emphasized technological self-reliance, digital infrastructure, and leadership in emerging technologies as central pillars of China's modernization strategy. More recently, AI governance and technological security have featured prominently in high-level Party discussions surrounding preparations for the next development cycle and broader debates on "high-quality development." During recent Politburo meetings and Central Committee discussions, Chinese leaders have repeatedly stressed the importance of developing "safe, reliable, and controllable" artificial intelligence while ensuring the "healthy and orderly development" of the sector. These discussions reflect Beijing's dual ambition: accelerating innovation and strategic competitiveness while maintaining political stability, social cohesion, and Party control over rapidly evolving technologies.

At the same time, China is seeking to shape international AI governance through initiatives such as the Global AI Governance Initiative, increased engagement in international standardization bodies, and the global expansion of Chinese digital infrastructure and AI technologies. These developments carry significant implications not only for China's domestic governance model, but also for Europe and the broader international community as they navigate questions of technological dependence, digital sovereignty, security, and regulatory coordination.

In this interview, Tom Abram discusses the evolution of China's domestic AI regulatory framework, the political and strategic logic underpinning Beijing's governance approach, and China's growing role in global AI governance initiatives. The conversation also explores the opportunities and limitations of EU-China cooperation on AI governance, highlighting both areas of potential convergence, such as technical safety and content labelling, and deeper normative divergences regarding state authority, individual rights, and technological sovereignty.

Conducted by Loke Sandvall, Katja Pomortseva, and Felicia Kämpe, this interview is divided into three parts from domestic to an international perspective. The exchange offers important insight into how China views AI not merely as a technological tool, but as a central component of economic modernization, political governance, and international influence.

China's domestic AI regulation

Loke Sandvall: The legal landscape of China's domestic AI regulation can be difficult to navigate. For non-experts, what are the key laws or policies one should understand to grasp China's national approach to AI regulation?

Tom Abram: What stands out is the growing concern in China, among both experts and authorities, around safe, controllable, and trustworthy AI, and the broad range of risks it entails. There is also a clear awareness of the need for governance at both the national and global levels. As early as 2018, Xi Jinping emphasized in a well-known speech on AI that China must ensure AI is safe, reliable, and controllable. Since then, this concern has gained momentum within the system.

To understand the governance of AI in China, a useful starting point may be: who regulates? I would distinguish three main actors, each reflecting a different aspect of governance. First, regulation stricto sensu: the Cyberspace Administration of China (CAC) is the central actor and has issued the most consequential rules so far. Second, ethics: the Ministry of Science and Technology (MOST) plays a role in both regulation and ethical governance. Its 2023 guidelines, for instance, require certain organizations to establish ethics review committees.

Third, standardization: the Standardization Administration of China (SAC) leads the development of national standards, notably through its technical committees on AI (SC 42) and cybersecurity (TC 260). Other bodies also contribute, such as the Ministry of Industry and Information Technology (MIIT) for industry standards.[1]

Sandvall: What are some key laws or policies?

Abram: China's domestic governance of AI has become quite developed. Contrary to the EU, China has adopted a more targeted approach with regulations addressing specific risks or applications, which is sometimes described as "vertical governance".

There are essentially five major regulatory instruments. The first, in 2021, focused on algorithms, particularly recommendation ones. It was followed in 2022 by a second provision on generated content, deepfakes more specifically, requiring clear labelling of synthetically generated material. Then in 2023, a broader regulatory framework on "generative AI services" was introduced, covering algorithms, content, and data in the context of the rise of chatbots. This aims to ensure that the training data is true and accurate and introduces licensing regimes for AI models.

The two most recent measures, both published in 2025, are particularly interesting. One requires most generated content to be labelled as such, both inscribed in metadata and with visible labels. Another draft targets AI companions and chatbots, calling on providers to strengthen the protection of minors and restrict harmful content. It was passed partly in response to public concern around anthropomorphic models, such as Xingye by MiniMax, and their social consequences.

Overall, China's domestic AI governance has

1 China differentiates different types of standards, notably national, local, industry, and military standards, each of them being drafted by different actors.

been very dynamic since 2021, addressing specific problems and risks as they emerge. At the same time, current provisions remain fragmented and incremental, which has led some Chinese experts to call for a more coherent and rationalised framework through a comprehensive national AI law. In 2023, the State Council – China's government – announced plans to draft such a national AI Law and submit it to the National People's Congress, though its final adoption may still take some time.

Sandvall: What factors have shaped the choice of starting with vertical regulation, and what advantages does this model offer?

Abram: Several factors may help explain this dynamic. First, there is a broader legal practice in China of regulating incrementally, starting with targeting very specific problems and only later trying to consolidate them into a more comprehensive and unified framework. This is something we have seen in other sectors as well, such as space and the regulation of commercial space activities.

Second, it may also reflect China's underlying objectives behind AI governance. The priority is not necessarily the protection of citizens and their individual rights per se, but first and foremost maintaining social stability and cohesion, preventing social incidents, and preserving Party control over information. This naturally drives authorities to focus on identifiable risks, emerging from specific applications like deepfakes or chatbots, that are seen as potential sources of disruption. In that sense, Beijing tends to opt for vertical governance even if, once again, it may not remain as such in a longer term.

The interesting aspects of this governance model are its responsiveness. It is quite effective at targeting concrete, pressing societal issues and regulating them early on, ideally before they scale. From the Chinese perspective, this is an effective way to manage a fast-moving technological environment in which new applications and risks emerge quickly. By contrast, the EU – until recently at least – tended to focus more on systemic and long-term risks. The Chinese approach, in turn, prioritizes adaptability, stability, and control in the face of rapid change.

Sandvall: Is manpower or population size a factor?

Abram: I don't think they frame it in those terms. Implementation is already complex, with enforcement carried out at multiple levels yet still centrally coordinated by national agencies beginning with the Cyberspace Administration. I don't think demographics is a decisive variable

Contrary to the EU, the Chinese have tried to put forward specific regulations to tackle specific problems. There are basically five major regulations and the last two are most interesting, both in 2025. One requires most generated content to be labeled, including visible labels and metadata. And another draft addresses AI companions and chatbots, pushing providers to strengthen the protection of minors and restrict harmful content.

in the equation. But the sheer scale of the AI ecosystem certainly is, as it makes oversight inherently challenging. Mechanisms such as the registration system for algorithms and AI models are designed precisely to map, monitor, and retain control over a large and rapidly evolving landscape.

Sandvall: How do policies balance innovation and risk? Is development prioritized?

Abram: They officially say they do not want to oppose regulation and development or innovation. They are trying to find a balance, which is somewhat still taking shape. However, AI is unambiguously perceived as a "core technology" and a strategic area for future economic development and international competition, particularly with the United States. Development remains the top priority, and authorities are careful not to constrain innovation and deployment too heavily. Programs such as "AI+" aim to spur concrete applications and integration of AI across economic and industrial sectors, an area where China has already demonstrated strong performance.

At the same time, not only has China put AI safety at the kernel of its AI governance approach as I have shown, but there is also a growing concern for the social implications of AI. The impact on employment for instance is a major issue actively discussed, including in bilateral or international settings. Labour was one of the important topics mentioned in the France-China joint declaration on AI in 2024. The awareness of such risks is closely tied to the Party's legitimacy and core interests in social stability and cohesion. As China's economic model evolves, AI is expected to play a major role, and the aim is to manage this transition in a way that maximizes opportunities while mitigating disruptions.

China's ambitions to shape global AI governance

Katja Pomortseva: In recent years, China has become more active in global AI governance debates. In 2025, it proposed a Global AI Governance Plan emphasizing international cooperation on using "AI for good" and the provision of open-source technologies to developing countries. What do these diplomatic efforts reveal about China's vision for reshaping global AI governance norms?

Abram: At the Belt and Road Forum for International Cooperation in 2023, China launched the Global AI Governance Initiative, which aligns with earlier initiatives beginning with the Global Development Initiative (GDI), but may also articulate with the Global Security Initiative (GSI) and Global Civilization Initiative (GCI). It is one of the broad shiny initiatives they have put forward, even though they have not provided that many details so far. But like the GDI or GSI, it provides an overarching framework through which China can structure and legitimize its subsequent actions in global AI governance.

In practice, China carries out a multi-pronged strategy aimed at positioning itself at the center of global debates on AI governance. One dimension consists in initiating high-level meetings on the topic. It is in those terms that we must understand the Chinese authorities' unprecedented support to the World AI Conference in Shanghai. Beijing really wants to make it a major global event for

AI governance. The 2024 edition marked a clear step up, with high-level political participation – Premier Li Qiang delivered the opening speech –, remarkable domestic and international coverage, and the adoption of the Shanghai Declaration, one year after the Bletchley Declaration.

Another key dimension is the promotion of inclusivity, particularly toward the "Global South" – a term that China has, by the way, started to use only recently in gradual replacement of "developing countries". At the Shanghai World AI Conference, China proposed the creation of an international AI governance body that would be located in Shanghai, especially championing the participation of Global South countries.

China also leverages existing multilateral and minilateral platforms. Officially, it advocates for UN-centered AI governance. Championing the inclusion of the Global South does indeed go through the UN arena as well. In 2024, for example, China and Zambia co-launched the Group of Friends on AI capacity-building at the UN, one of many informal coalitions fostering dialogue on these topics. But Beijing is very, if not more, active outside of the UN system as well. BRICS happened to be one of the platforms for China to push discussions, most notably on AI ethics. In the 2025 meeting, BRICS have also put forward a statement on global governance of AI, linking AI to the Sustainable Development Goals and the need for responsible use.

These different initiatives go hand in hand with China's promotion of some key notions to shape the global discourse on AI governance, such as "AI for good," "people-centered approach" to AI governance, and the overall framework of "building a community of shared future for mankind in the AI domain".

Standardization is another aspect we tend to overlook. China is highly active in international standard-setting bodies, promoting its own standards and norms. Interestingly, these are also one of the last multilateral fora where the U.S. still maintains active participation.

Taken together, these efforts reflect several objectives. First, shaping global norms in ways that remain compatible with regime security and core political interests, including limiting the scope of human rights-based constraints. Second, presenting China as a responsible power in both AI development and global governance, and a champion of multilateralism and South-South cooperation. Finally, consolidating its status as a major power – and more specifically a major science and tech power – capable of shaping global debates.

Yet, there is a paradox. Despite substantial state and, perhaps more importantly, non-state

China's efforts reflect several objectives. First, shaping global norms in ways that remain compatible with regime security and core political interests. Second, presenting China as a responsible power in both AI development and global governance, and a champion of multilateralism and South-South cooperation. Finally, consolidating its status as a major power – and more specifically a major science and tech power – capable of shaping global debates.

resources and expertise, Beijing still struggles to fully engage with and exert influence in some of the main arenas of global AI governance.

Pomortseva: Through initiatives like the Digital Silk Road, China offers technical assistance and data security systems to developing countries. What implications might this have for Western efforts to promote shared standards for ethical AI?

Abram: As with other digital technologies, this raises a familiar issue: the export of services and tools through initiatives like the Digital Silk Road, not only to the Global South but also, to some extent, to Europe. One of the key questions is what these technologies carry in terms of values, usage, norms, and economic influence. This is partly what some scholars refer to as the "Beijing effect," highlighting the role of Chinese tech firms in diffusing norms and shaping transnational data governance through market expansion.

In the AI field, this is for instance visible with large language model providers seeking to expand abroad, often offering relatively low-cost solutions that are particularly attractive in "Global South" countries. Some smaller startups like the "four tigers" sought to expand their market out of China and sell solutions to African countries for instance. While these efforts may be indirectly supported by the state, they are also, and I would argue primarily, driven by pure commercial interests. At the same time, they align with broader policy frameworks like the Digital Silk Road, which remains a central component of Xi Jinping's agenda.

This connects directly to the issue of standards. Diffusion of technology is also a way to diffuse norms: the more widely Chinese systems are adopted, the greater the ability to shape technical standards and, over time, governance practices. This has clear implications for Western efforts to promote shared approaches to ethical or responsible AI.

That said, and contrary to what we sometimes hear, Chinese actors are also truly engaged in ethical discussions and express a willingness to increase international dialogue on ethics. While there are clear differences in approach, I am convinced there are a lot of commonalities. This suggests that, alongside competition over standards, there remains space – and arguably a need – for deeper engagement and dialogue on AI ethics, through existing frameworks such as those led by UNESCO, and beyond.

Pomortseva: As China emerges as a pivotal supplier of AI infrastructure and foundational models, while the EU remains structurally

Chinese actors are also truly engaged in ethical discussions and express a willingness to increase international dialogue on ethics. While there are clear differences in approach, I am convinced there are a lot of commonalities. This suggests that, alongside competition over standards, there remains space – and arguably a need – for deeper engagement and dialogue on AI ethics, through existing frameworks such as those led by UNESCO, and beyond.

dependent on extraterritorial technologies, to what extent does this asymmetry constrain Europe's strategic autonomy in negotiating AI governance norms with Beijing?

Abram: Our Europe's dependence on extraterritorial technologies across hardware, software, models, and critical components is obviously a major structural constraint, and one that EU institutions have become increasingly aware of. This dependence is twofold: on the United States, but also on China in certain segments of the value chain. It inevitably shapes Europe's room for maneuver, both in domestic regulatory choices and in its external positioning on AI governance.

This is all the more problematic given the divergences between the EU and the U.S. on data and digital governance. Transatlantic tensions in these areas are well documented, and they illustrate how technological dependence can feed back into regulatory decisions. The EU's data regulations have recently faced repeated assaults from the U.S. administration. Recent debates and decisions by the EU Commission to adjust the implementation timeline of the EU AI Act, partly framed in terms of competitiveness, also reflect these underlying constraints of dependencies.

The constraint stemming from our dependence on the U.S. and China is not absolute. On the one hand, we should be cautious not to overstate the so-called "Brussels Effect," which has interestingly evolved from an analytical concept into a more politicized narrative. I have sometimes observed a tendency among Brussels' elites to rely on past successes, assuming that EU regulations such as the AI Act will automatically set global standards, at times before the Act was even finalized. That said, the EU still retains significant regulatory leverage through the size of its market and its ability to impose rules that external actors may need to accommodate. The key challenge, therefore, is to balance this regulatory influence with underlying technological dependencies.

EU-China cooperation on AI

Felicia Kämpe: **Some scholars have called for a high-level EU-China dialogue on AI governance. From your perspective, should cooperation be deepened in this area? If so, what benefits and drawbacks would this entail?**

Tom Abram: If I may take a French perspective to answer, the question you raise touches upon the broader debate on where we should cooperate with China, in a context of major risks resulting from economic interdependence and trade imbalance. As Macron said, balancing economic relations with China is a question of "life or death"[2] for European industries.

Against that backdrop, the key question is not whether to cooperate, but where and how. Climate change has traditionally been a relatively consensual area of engagement. My view is that AI governance should be another. There are clear reasons to deepen dialogue, both for the overall EU–China relationship and for the broader objective of governing AI and retaining some degree of control over its safety and societal impacts.

2 This is a quotation attributed to French President Emmanuel Macron, as reported by Tom Nicholson in Politico, "Europe must rethink China strategy, Macron says — and China needs to help," December 7. 2025, https://www.politico.eu/article/europe-china-emmanuel-macron-foreign-investment-trade/

There is a perceived gap in values and AI use. In Europe, the Chinese model is often seen as a cautionary example, associated with the use of AI for surveillance and social control. The EU AI Act explicitly restricts practices such as facial recognition in public spaces and social scoring, tools that are more widely deployed in China. For this reason, Europe has, in recent years, been keener to cooperate with so-called "like-minded" countries like Japan or Canada.

Both the EU and China are, in different ways, leading actors in AI regulation, and Beijing has shown interest in engaging with European approaches, including the AI Act. This creates a basis for exchange, and potentially for convergence on certain issues, particularly where risks are transnational. Beyond the inherent technical risks of AI models and data, such as bias or hallucinations, and broader social impacts like employment, misuse of AI should be a global concern. AI-driven or enabled misinformation, cyber-attacks, and other threats to citizens and national security are significant and inherently cross-border. Precisely for that reason, they require some level of international coordination.

Needless to say, cooperation should not imply naivety either. The challenge is to pursue targeted and pragmatic cooperation, engaging where interests align, particularly on risk mitigation and safety, while remaining attentive to discrepancies and potential vulnerabilities.

Kämpe: Despite differing normative priorities, such as the EU's emphasis on individual rights and China's focus on state sovereignty, both sides share concerns about technical safety and reliability. In which areas do you see the greatest potential for practical EU-China cooperation?

Abram: There are several areas where dialogue and convergence is possible. On some issues, the EU could also draw lessons from the Chinese experience. This is notably the case for the regulation of deepfakes and content labeling, where China has moved early, as well as for governance tools such as algorithm registries.

China's registry system for AI services is particularly interesting in this regard. It distinguishes between algorithms used for online information services and those used for content generation (AIGC). Registry requirements target especially those systems deemed capable of "shaping public opinion" or "mobilizing society." In practice, this translates into four major registration lists for algorithms and AI models managed by the Cyberspace Administration, depending on their functions and nature. These mechanisms allow for a certain degree of oversight, monitoring, and evaluation of AI systems deployed in the country.

Such tools could provide inspiration for European policymakers, even if they would not be transposed in the same form, given different legal and normative frameworks.

Kämpe: What are the major obstacles hindering cooperation?

Abram: Several obstacles stand out, most of them rooted in diverging approaches to governance and underlying normative differences.

First, there is a perceived gap in values and AI use. In Europe, the Chinese model is often seen as a cautionary example, associated with the use of AI for surveillance and social control. The EU AI Act explicitly restricts practices such as facial recognition in public spaces and social scoring, tools that are more widely deployed in China. For this reason, Europe has, in recent years, been keener to cooperate with so-called "like-minded" countries like Japan or Canada.

Second, economic and strategic competition remains a major obstacle. AI is largely framed through a competitive, economic lens, with China and the United States viewed less as partners than as technological rivals. This logic of competition and even of a "race" complicates cooperation on governance. Some European officials and experts even argue that both China and the U.S. seek to preserve their advantage and limit the EU's emergence as a global AI power, including through efforts to impose global regulatory frameworks.

Third, there are internal divergences within Europe itself on the very need to regulate. Some member states, including France and Germany, have expressed concerns about overregulation and its impact on competitiveness. Emmanuel Macron, for instance, has argued that Europe "is overregulating while underinvesting", thereby missing economic opportunities. Others just do not believe in the need to regulate, and downplay risks related to AI safety or misuses.

More fundamentally, there is a structural difference in governance philosophy, which I would summarize as "rule of law" versus "rule by law." The EU's approach is centered on protecting society and citizens' fundamental rights from potential abuses primarily by state actors, aligned with its tradition of the rule of law. In China, the emphasis is different: regulation is designed to legally empower the state to protect society, and more specifically China's social and political order, from AI-related risks and misuses by non-state actors. I am not saying that one is better than the other, but that the EU and China took the problem from two opposite angles, which may limit convergence.

Kämpe: **Finally, is there anything else you would like to add?**

Abram: We haven't touched upon the governance of military AI more specifically. China is integrating AI into its military in many ways – just as the US and EU member states are. This remains a complex and sensitive issue, but dialogues do exist, whether on military affairs more broadly or on the impact of AI in specific domains such as biosecurity or nuclear deterrence. We need to deepen these discussions and work toward identifying lowest common denominators, creating confidence-building mechanisms, shared understandings of risks, and best practices, which could eventually support more binding forms of regulation. The conclusion I draw from my participation in some of these dialogues is that there are real opportunities for convergence, despite increasingly intense geopolitical competition.

Afterword

Maud Descamps

The interview highlights the complexity and dynamism of China's approach to AI governance. Rather than pursuing a single comprehensive framework from the outset, Beijing has adopted a highly adaptive and sector-specific regulatory model, seeking to address emerging risks while simultaneously supporting rapid technological development. This reflects broader features of China's governance model, where regulation is closely linked to questions of social stability, political control, and strategic competition.

At the international level, China's growing activism in AI governance demonstrates that the contest surrounding AI is not limited to technological capabilities alone, but increasingly concerns the power to shape global norms, standards, and institutional frameworks. Through initiatives such as the Global AI Governance Initiative, the Digital Silk Road, and active participation in international standardization bodies, China is positioning itself as a central actor in defining the future governance architecture of AI.

For the European Union and its members and allies, the discussion underscores both opportunities and challenges. While important differences remain regarding human rights, surveillance, and the role of the state, there are also areas where practical cooperation may prove both necessary and beneficial—particularly in fields such as AI safety, content authenticity, technical standards, and the governance of emerging risks. At the same time, growing technological dependencies and strategic competition complicate the EU's pursuit of digital sovereignty and autonomous governance capacity.

Ultimately, the interview suggests that AI governance will increasingly become a key arena of geopolitical negotiation. Navigating this landscape will require not only technological competitiveness, but also sustained dialogue, institutional innovation, and international cooperation capable of addressing the profound societal and security implications of AI development, including in sensitive domains such as military and nuclear applications.