China's AI-Powered Surveillance State

Valentin Weber

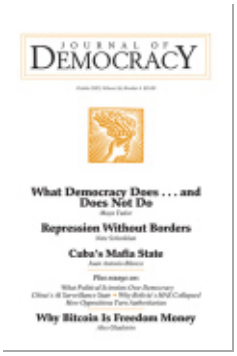

Journal of Democracy, Volume 36, Number 4, October 2025, pp. 151-160 (Article)

Published by Johns Hopkins University Press
DOI: *https://doi.org/10.1353/jod.2025.a970356*

CHINA'S AI-POWERED SURVEILLANCE STATE

Valentin Weber

Valentin Weber *is a senior research fellow at the German Council on Foreign Relations and a China Foresight associate at LSE IDEAS. He has been an Open Technology Fund Senior Fellow in Information Controls at Harvard University's Berkman Klein Center for Internet and Society, and holds a doctorate in cybersecurity from the University of Oxford.*

Artificial intelligence (AI) is transforming not only economies, but how law enforcement and surveillance are conducted. The People's Republic of China (PRC) has rolled out the largest AI-based public-surveillance system in the world and handed over to AI ever more of the tasks that human police officers once handled. Like drone pilots, police in Chinese cities increasingly sit in command centers looking at screens that allow them to monitor and manage public security using data aggregated, sifted, and visualized by AI. Human officers would command and coordinate police units based on such data, but improved large language models (LLMs) are taking over this "back-end" role. On the streets, Chinese advances in robotics mean that AI-governed robots will move from their current role supporting human security personnel and start directly arresting dissidents.

On the tactical level, AI is increasingly able to run the whole law-enforcement cycle: gathering information, deploying units, commanding, planning, and patrolling. On the strategic level, AI's impact on the PRC surveillance state is even more profound. Technology has made massive data integration possible not only within cities, but on the level of whole provinces and indeed the entire country (rural and wild areas included). The Chinese Communist Party (CCP) will soon be setting up a national digital command center. The current PRC surveillance state is corrupt, costly, and takes a lot of human effort to run. The CCP wants one that will require little human participation, being "staffed" instead by AI working through drones, autonomous vehicles, and humanoid robots.

Will this make the communist party-state that rules China forever

Journal of Democracy Volume 36, Number 4 October 2025

dominant? Will that regime ever fully control human emotions or eliminate human unpredictability? While the CCP can reduce its reliance on human police officers, it will remain stuck with vast numbers of humans living within the PRC's borders. These people will remain highly inventive and unpredictable, especially when facing economic or personal distress. If the CCP falters—say by failing to deliver economic prosperity—then all the robots, drones, and cameras in the world will not be able to contain the unrest.

Twenty-five years ago, the PRC began digitizing its ministries, including the police and public services. More surveillance cameras were mounted. Today, there are seven-hundred million—a camera for every two citizens. At the same time, there are more mobile phones in China than actual citizens. Surveillance devices—whether phones, smart meters, smart watches, or other pieces of linked technology—already greatly outnumber the police officers, paramilitary troops, and informants who maintain regime security for the CCP leadership. Online devices are the eyes and ears of the party-state, its key to situational awareness.

For the past fifteen years, China's focus has been on data integration. Cities and provinces have used different firms to supply technology, which means that systems are not always interoperable. Even within cities, the police often use software that does not work with other city services. Cities were the first focus of data-integration efforts, which now extend to provinces. The Beijing-Tianjin-Hebei urban conurbation, home to around 109 million people, saw the first cross-provincial work on sharing "big data," and Jiangsu-Shanghai-Zhejiang (174 million people) has been added to the list as well. The CCP can thus monitor more than 20 percent of the PRC's 1.4 billion people via a single pair of screens. The party-state's dream of "one-screen governance" draws closer to realization.

The CCP's higher-ups are eager to get direct, ground-level views of unfolding events everywhere in China, with no need to rely on reports filtering up through layers of bureaucrats with their own agendas. Local officials will lose both their ability to soft-pedal riots and mishaps, and their capacity to overstate successes. Cameras, computers, and robots do not angle to avoid punishment or gain promotion (at least not yet), so the CCP top ranks feel confident that they are receiving a clearer picture of "the real China" than ever before.

As the geographic reach of surveillance technologies grows, so does the CCP's control. The skies above Chinese cities swarm with drones. Through them, securocrats can monitor beyond city limits too. Countryside surveillance is becoming more common. Futility will be the result of any rural escape such as the one that *Brave New World*'s protagonist John Savage tried by fleeing to the lighthouse in Aldous Huxley's 1932 novel. In 2020s China, the surveillance state haunts the rustic greenery as well as the urban asphalt. Hangzhou, the hometown of DeepSeek and

the PRC's most technologically advanced city, could track cars a decade ago and now tracks pedestrians, cyclists, and the flow of goods through airports and railway stations.

Another advantage that PRC authorities hope to glean from AI is the ability to react to citizens' queries and complaints before people become frustrated, angry, and unpredictable. AI with reasoning capabilities is being integrated into municipal services. Chatbots answer citizens' questions about administrative matters. In Hangzhou, City Brain GPT offers citizens digitized "civil servants."[1] The latest AI, especially DeepSeek LLMs, can enable police to search hundreds of thousands of surveillance cameras in real time to find, say, a person in a green hoodie or a particular car doing anything unusual over the past week. In addition, AI can suggest to the police where to deploy patrols and at what times, or how many officers might need to be sent to a particular crime or accident scene.

In 2020s China, the surveillance state haunts the rustic greenery as well as the urban asphalt. Hangzhou, the hometown of DeepSeek and the PRC's most technologically advanced city, could track cars a decade ago and now tracks pedestrians, cyclists, and the flow of goods through airports and railway stations.

DeepSeek's newest reasoning LLM is cheap enough to use that it will support the broad rollout of AI "agents" (decisionmaking software). This will shift surveillance from passive (cameras and sensors see and hear) to active mode. A person might have an AI agent make travel reservations, but if the CCP asks the AI agent to limit a dissident's geographic movements, then those reservations will be canceled. The police could instruct all AI agents to make sure that a certain person never leaves her neighborhood. Or an agent could disable the purchase of any personal-transport services via a person's WeChat Pay account without blocking that same account's ability to pay for food or routine bills.

In March 2025, Lenovo introduced the Urban Super Intelligent System as an upgrade of the City Brain city-management platform.[2] The new system can execute decisions, not merely suggest them. Lenovo's product features a single AI superagent that coordinates the AI agents carrying out decisions. Commercial and transport AI agents, for example, exchange information regarding traffic, sales data, passenger flows, and consumer preferences. Public-security AI agents would work together with commercial and other AI agents to execute tasks autonomously.

After a protest demonstration, in a hypothetical scenario, a human police officer might instruct the AI superagent to make sure such a protest will not recur. The AI superagent would then ask AI security field agents to identify everyone involved, call up their movement patterns,

and block them from using public transit for two days while also keeping all their personal vehicles inside geographic limits. A propaganda AI agent would create news and social-media content putting the protests in a bad light while sending individual private messages to the entire social networks of all those seen protesting, to mark them as people whom friends, relatives, and acquaintances should avoid.

In order to "rehabilitate" the protesters, the AI superagent would command the city's public-services AI agent to have every demonstrator spend several hours a week doing community service and an hour a day interacting via phone or computer with a recovery AI agent that would mix interrogation questions with ideological indoctrination. The AI superagent would keep tabs on ideological progress and adjust instructions to AI field agents accordingly. Lack of progress or refusal to participate would mean punitive measures staying in place or even growing harsher. Human officials would need to become involved only if anomalies appeared.

Lenovo selected Wuyishan and Yichang as the first two cities to take this next step of smart-city development using AI agents. Yichang boasts extensive hydropower (and hence cheap energy for computing), while Wuyishan is a bustling tourist center. The human role in managing those two cities is receding toward shallow participation.[3] For now, humans are setting goals and approving crucial decisions, but the system may soon no longer need them. Increasingly, machines talk to other machines to coordinate all phases of surveillance and control.

How AI Fills the CCP's Security Gaps

One major vulnerability of the surveillance state is that it is still disembodied. As Minxin Pei writes:

> Over the past eight decades, the CCP has constructed a vast network of millions of informers and spies whose often unpaid work has been critical to the regime's survival. It is these men and women, more than cameras or artificial intelligence, that have allowed Beijing to suppress dissent. Without a network of this size, the system could not function.[4]

It is police officers who do regular inspections to ensure that internet cafés are recording people's identification cards properly, or who check on key individuals (Tibetans, Uyghurs, dissidents). Pei is correct that intimidation works best if a person knocks on the door, and he is also right in pointing out that the PRC's surveillance infrastructure is still mostly passive. Mobile-phone location and facial-recognition cameras can tell the CCP who was at an internet café and for how long, but robots do not yet conduct "in-person" checks on key individuals.

During the past decade, the CCP has had eyes and ears, and even a brain to process and make sense of data. It has 5G networks and fiber-

optic cables that serve as neural pathways carrying information from closed-circuit television (CCTV) cameras to cloud-computing centers. But it all remains passive. And even when the surveillance state becomes active via AI agents, those actions remain mostly on the internet. The surveillance infrastructure can only block people's access to services. People's train tickets might not work. Their connected vehicles may not start. Their emails and texts may be screened by censorship software.

The technical surveillance infrastructure has lacked limbs and thus has needed to rely, as Pei notes, on human enforcers to put "boots on the ground." The CCP has been seeking to change this. The first step was to introduce aerial drones that could follow protesters around a city. Autonomous vehicles fulfill a similar purpose. The increasing capabilities of autonomous cars means that they will soon be able to do more than be remotely ordered not to start. What if a dissident got into her car one day and found the doors were being locked remotely (with the locks then rendered inoperable by anyone inside, also by remote command) and her self-driving car driving itself to a police station with her trapped inside? This kind of dystopian nightmare scenario is the CCP's dream.

Drones and self-driving cars are not all the CCP wants, however. To put enforcement by "embodied AI" on the sidewalks and into buildings, the party-state wants to deploy humanoid robots run by AI and with physical limbs whose abilities will soon rival or surpass those of humans. In March 2025, the Binjiang Public Security Bureau in Hangzhou introduced a humanoid robot called "Bin Xiaoxin."[5] It belongs to a network of aerial drones, two- and four-legged robots, and driverless automobiles. The machines have facial-recognition capabilities and can create three-dimensional virtual snapshots of crime scenes. As machines, they will need downtime for recharging and maintenance, but the idea is to have enough to make full-spectrum patrolling possible around the clock.

Pei argues that the CCP surveillance state is labor-intensive in part because the party-state wants to avoid having to rely on a single agency and stream of information and so has created a number of different security and intelligence organs. This gives the highest authorities a more rounded picture of what is going on throughout China while also fostering bureaucratic rivalries and "turf wars" that can be manipulated from above to keep the security apparatus as a whole in check and subject to the party. The CCP's answer to the question "Who watches the watchdog?" in other words, is to have several rival watchdogs. An implication of this is that the secret police are not the central authorities' only source of information.

A parallel stream of information is gathered by CCTV cameras and other sensors, analyzed by AI, and then visualized on large screens in city and cross-provincial command centers for viewing by senior securocrats—all within moments. CCP chairman Xi Jinping and his colleagues

are glad to have information relevant to public security that AI has aggregated and analyzed, and that bureaucrats below have had no chance to meddle with, shape, or spin. China has had an "emperor" for many centuries, but never one with the panopticon in whose all-seeing center Chairman Xi sits.

There are still some geographic limits, so the panopticon remains a work in progress. Hebei province neighbors Beijing and shares an AI "cross-provincial brain" with the national capital. Information from distant Tibet, by contrast, still has to travel through layers of bureaucrats. Before long, however, data from the streets of Lhasa, Tibet's capital, might be flowing straight to Beijing in real time.

The new security structure's "edge devices"—drones, robots, cameras—have small digital brains that fulfill very narrow intelligent goals, such as picking a fugitive's face out of a crowd. This information is sent to a city brain, yet another layer of intelligence, where the larger goal of preventing widespread unrest is pursued. In short, the Chinese surveillance state functions like an octopus. Each of the eight arms (brains) can act independently, but the central brain is still able to apply top-down control. Wu Zhiqiang of the Chinese Academy of Engineering views it similarly. He sees the city as a system comprising a main brain and auxiliary brains. Together, they create "multibrain" social intelligence, which learns how social communities function and uses the collaboration of many digital brains to govern them.[6]

The CCP's Remaining Vulnerabilities

The CCP's growing reliance on advanced technology brings a new vulnerability: rogue AI. There are already LLMs that lie, disable their oversight mechanisms, or make their way onto external servers. As AI becomes more powerful, this tendency might increase. Public-security AI in China is designed to detect complex patterns, gather real-world experience, and learn from social interactions. Each humanoid robot is constantly adding to its own set of experiences. AI trains itself on what it takes in. Embodied AI agents are designed to pursue autonomous learning and self-improvement. If Robot A piles up experiences that lead it to become good at pursuing people, this could eventually single it out as a proficient hunter.

The current focus of developers in China is to make robots more like humans, to make them learn as fast as toddlers, using real-world data and not merely what is available on the internet. As robot police officers learn to become more humanlike, could they learn (among other things) how to "cut corners" in pursuit of the basic goals that have been programmed into them? This would be a species of robotic corruption, ignoring rules not for bribes, but for the sake of greater efficiency. If a patrol robot were to perceive instructions from a human police offi-

cer as demanding inefficiency, could the robot become prone to disable its oversight mechanism? There might be times when human officers would not want their "proficient hunter" robot to pursue as aggressively as it can, but will the robot always listen? And that scenario involves only one small digital brain in a single edge device. What if something like this happened on a system level to the multibrain of an entire city or province?

For now, robots are only allowed to identify suspicious behavior, they cannot arrest people, but this could change. They are authorized to guide people but cannot appease people who are angry or emotional. The PRC police follow a maxim of "digital assistants, not digital deciders."

Chinese police officers know what AI-powered robots can do, and maintain that humans are needed for their emotional intelligence and their ability to manage complex crises. But what if machines are or soon will be better at these roles? LLM chatbots are already serving some humans as therapists, life advisors, and even romantic partners. Similarly, robot police officers are good at listening and reflecting while not every human police officer has great skills or inclinations in these areas (to say nothing of human officers who are outright brutal or corrupt).

What about crisis management? Are humans untouchable on that score? Robots are already used in complex operations that carry high risks of harm to human police officers: Better to have robots defuse bombs or burst into violent criminals' suspected hideouts in the middle of the night. As drones and humanoid robots become more capable, their use instead of humans to handle crises will only increase. What is most striking is the change in how PRC police officers see themselves. AI has made them technology managers, risk predictors, and social mediators, and they know it. Has the technology become so fundamental to their work that they rely on it too much?

Even without "going rogue," AI can threaten stability. Constant reliance on it could lead human police to lean on it and trust it too much. As systems become more complex, they will become more opaque. Will this cause securocrats to ask if they should entrust public security to a system that they do not fully understand?

A subtler problem for the CCP regime is the rising level of citizen expectations that might flow from the elimination of both police corruption and multiple layers of party-state bureaucracy. Citizens might well ask: If the center has immediate control and oversight, why do our grievances seem to go unheard? The PRC's security agencies have long tolerated complaints about local officials as a way of deflecting discontent away from central CCP leaders. A future national digital brain in Beijing will imply that central officials know about and are responsible for what goes on at the local and provincial levels.

Strikingly, the CCP has gone to great lengths to keep physical dis-

tance between rulers and ruled. Restricting dissidents' movements is common, with phone tracking used to ensure that they do not reach a part of a city in which government offices or other sensitive regime properties are housed. This focus on physical limits is in keeping with the CCP's rating of "in real life" protest as a graver threat than online dissent. Chinese securocrats feel that they have perfected online control but still struggle to find the handle offline. Their main goal is to prevent large crowds from gathering. In 2022 and again in 2024, they failed notably at this.

On 13 October 2022, banners appeared on Beijing's Sitong Bridge demanding freedom, elections, an end to lies and covid-19 lockdowns, and the removal from power of "dictator" and "national traitor" Xi Jinping. The banner hanger, named by some as Peng Lifa, is said to be still in detention. In November, anti-lockdown protests began breaking out. Thousands took to the streets to vent their anger against the PRC's movement restrictions, which enforced Xi's "zero-covid" policy and were some of the strictest in the world. The demonstrations are thought to be the largest antigovernment manifestations since the Tiananmen Square movement that the CCP violently crushed in June 1989. Unlike in that earlier case, they were not limited to one city.

These protests revealed the CCP's greatest vulnerability: Its own incompetence. The virus at the root of the trouble may have been of the PRC government's own making, seems 80 to 90 percent likely to have escaped the laboratory in Wuhan due to bungling (says Germany's Federal Intelligence Service), and was met by harsh, unpopular, and misguided CCP-ordered measures in pursuit of "zero covid" that stirred public anger while failing to contain viral spread.[7]

Most remarkably, the largest protests occurred in the giant coastal city of Shanghai, where local people turned out to support ten fellow citizens who had died more than three-thousand kilometers away in a November 24 high-rise fire in Ürümqi, the capital of the Xinjiang Autonomous Region. Lockdown measures had been at least partly to blame for the deaths, protesters believed. They held up sheets of blank white paper to signify mourning and their opposition to government policies.

The Midnight Cyclists of Kaifeng

Two years later, tens of thousands of students in Henan province took up bicycling the 65 kilometers from Zhengzhou to Kaifeng. These after-dark "dumpling runs" were a campus craze. Smiling young people pedaled their way through two-wheeling all-nighters, flying the PRC flag from their rideshare bikes and spending their breakfast money in Kaifeng, a picturesque and historic town that is one of China's "Eight Ancient Capitals." Authorities had their AI, their drones, and their CCTV cameras to monitor the jaunts for weeks, but were nonetheless shocked

by how popular the rides became, abruptly shutting them down in early November 2024 as cyclists clogged the highway and created safety concerns. Despite all the high-tech tools at their disposal, officials had been surprised by a random outburst of pure human whimsy backed by the energy of youth and the power of "going viral" on the internet.

As long as humans remain in control of the surveillance state, as long as they make decisions, the security of the communist party-state will remain in doubt—humans simply make too many mistakes for this not to be the case. Hence the CCP leadership's dilemma: Should it entrust security decisions to AI systems that no one may fully understand, and that might go rogue or run errors? Or should the reliance on human judgment be retained even though an error at the wrong place and time could cost the regime its life?

Soviet communism lasted 74 years, from 1917 to 1991. Its Chinese counterpart marks its seventy-sixth anniversary in October 2025. Today, digitized repression is a major advantage the PRC has that the USSR did not. Research has shown that regimes which rely on digital repression last longer than those that do not. In communist East Germany, the secret police gathered more data than their analysts could process, a task now made trivially easy by electronic records and AI. Nobody does digital surveillance and control better than the PRC. It is the world's leader in these, and looks to remain so for some time.

The CCP will keep fixing its vulnerabilities. People will be ever more closely monitored, and their ability to move about will be subject to ever more fine-tuned tightenings and loosenings based on ever more detailed AI assessments. The CCP will maintain AI repression; Xi Jinping is well aware of how the USSR fell and is determined not to let the Chinese party-state go down that path. He focuses on the economy as his immediate key to legitimacy, but does not forget to keep tight clamps on freedom of speech and association, liberty of thought and conscience, and dissent.

Even so, the CCP regime's security will remain profoundly vulnerable. Human nature is messy, irrational, and unpredictable—with manifestations not always as innocuous as some lighthearted bike rides. Those were harmless in themselves (aside from traffic tie-ups), but they exposed the CCP's Achilles' heel. Authorities made blundering judgments, letting the cyclists slide only to panic-ban them once the fad went viral and the bike fleets got too big.

Poor decisionmaking—including at the highest levels—is unlikely to go away. Who would have thought that deciding to breed highly contagious viruses in a Biosafety Level 4 lab (officially the safest kind) in Wuhan could lead to fear for survival among leaders in Beijing? Or that the zero-covid policy would be so severely unpopular?

Under conditions of liberty, human unpredictability can go from being a weakness to being a strength. Where civil society is free, where people can talk and gather in peace, and where the political system is

competitive, public discontent and desire for change can be treated as normal aspects of the system that the system can accommodate by means of democratic responsiveness. Thus, publicly aired grievances and open protests need not be seen as purely alarming shocks threatening to topple a brittle authoritarian edifice. Instead, and within broad limits, grievances and protests can provide useful democratic feedback that can be weighed and (after open debate and discussion) perhaps acted upon by freely chosen representatives of the people. Democracy, in short, has a built-in resilience that autocracy lacks.

The anger expressed during the covid protests and the joy shared by thousands of people biking peacefully between two cities were two very different emotions, but both underlined the vulnerability of the CCP. Human emotions and unpredictability will remain the party-state's constant headache, especially when its functionaries mess up, as they have in the past and surely will in the future. No amount of AI surveillance and control is going to change that dynamic, and tech cannot make a brittle system resilient—only democracy can.

NOTES

1. Zhao Lu and Tu Youju, "Hangzhou Fully Launches the Construction of 'City Brain 3.0,'" *Zhejiang Daily,* 31 March 2025, *https://baijiahao.baidu.com/s?id=1828119933477355350*.

2. "Another Success! The World's First Batch of Urban Super-Intelligent Bodies Launched in Yichang," *China Daily,* 20 March 2025, *https://baijiahao.baidu.com/s?id=1827109397143003768*.

3. Zeng Xiangling, "'Urban Super Intelligence' Is Launched: Lenovo Smart City 4.0 'Comes to Fruition,'" Baidu, 25 March 2025, *https://baike.baidu.com/tashuo/browse/content?id=82a3cad5d1ba2514648fe7df&lemmaId=65504551&fromLemmaModule=pcBottom&lemmaTitle=%E5%9F%8E%E5%B8%82%E8%B6%85%E7%BA%A7%E6%99%BA%E8%83%BD%E4%BD%93&fromModule=lemma_bottom-tashuo-article*.

4. Minxin Pei, "Why China Can't Export Its Model of Surveillance," *Foreign Affairs,* 6 February 2024, *www.foreignaffairs.com/china/why-china-cant-export-its-model-surveillance*.

5. Lu Tianran, "When Robots Wear Police Uniforms: Can the New 'Police Officers' on Hangzhou's Streets Break the Police Iron Rice Bowl?" Baidu, 24 March 2025, *https://baijiahao.baidu.com/s?id=1827349112869115832*.

6. "To Build a True City-Level Brain, AI Needs to Master 'Group Wisdom' and Learn to Play Games," *Shanghai Observer,* 8 July 2023, *https://web.archive.org/web/20250423075321/https://export.shobserver.com/baijiahao/html/630479.html*.

7. Francesca Gillett, "German Spy Agency 'Believed Covid Likely Started in Lab,'" BBC, 13 March 2025, *www.bbc.co.uk/news/articles/cz7vypq31z7o*.