

**New York City Council
Committee of the Whole**

Reasserting AI Safety for New Yorkers
October 5, 2026

Written Testimony of Aashna Agarwal
Program Associate, AI Now Institute

Thank you Speaker Menin and the members of the City Council. My name is Aashna Agarwal and I am the Program Associate at the AI Now Institute, an independent organization producing policy research on artificial intelligence. In particular, we focus on policy strategies to shape artificial intelligence in the public interest and redirect the AI industry away from its current path: unbridled commercial surveillance, power concentration in a few dominant firms, and a lack of public accountability. With those interests in mind, I am testifying today regarding the risks posed by AI.

Introduction

We are at a critical inflection moment in the trajectory of AI technology development and deployment, and in the policymaking efforts to safeguard the public from its harms. The companies building Artificial Intelligence have never been louder about “safety.” They speak in abstractions—warning about the “existential risks” that come with great technological advances and for which only they have the solutions. But safety is not an abstraction that lives at the frontier. For New Yorkers, safety materializes in how AI is embedded in the systems we use and rely on every day, from our hospitals to our utilities, from our city services to our workplaces, often to the detriment of the people those institutions exist to serve.

My testimony today addresses three critical ways that AI companies’ capital interests conflict with safeguarding the public from the risks posed by the very technology they are developing, as well as how the Council can help New York City resist the integration of harmful technology and protect the safety of New Yorkers.

For one, we need to **maximize friction when it comes to the integration of AI into sensitive social domains**. Second, we need to **remedy and redress the harms AI has brought to New Yorkers already**. Third, we need to **resist the concentration of power in the AI industry and the single points of failure it creates, which put public safety and the resilience of City infrastructure at risk**.

1. The Rapid and Risk-laden Integration of AI into New York’s Most Sensitive Systems

The scale of AI companies’ capital expenditures is staggering. Goldman Sachs projects AI-related investment to reach \$1 trillion globally this year, including \$581 billion of spending in the U.S.¹ Yet these stratospheric expenditures raise a question the industry has yet to answer convincingly: where will the revenue come from? By one estimate, the industry would need to generate over \$6 trillion in annual revenue by 2031 to justify the spend, and is on track to fall nearly \$4.2 trillion short of that mark.²

The AI gold rush has reached a fever pitch. And, like every gold rush before it, the surest money isn’t in striking gold but in selling shovels. AI companies know this and have accordingly sought to sell their technologies into public institutions that customers cannot easily walk away from. Companies push their products into every corner of public and private life as quickly as possible—into our hospitals, our utilities, our workplaces, and our city agencies—whether or not those products work as promised. These are systems that New Yorkers can least afford to see fail.

Where companies can’t sell directly, they leverage their deep pockets to offer cheap and discounted services as a way to lock in future customers. In August of last year, OpenAI and Anthropic offered their models to federal agencies for \$1 per agency for a year, while Google offered Gemini for 47 cents.³ These trial periods were attempts to lock agencies into paid products that will eventually prove too costly or disruptive to give up. This playbook is not new. In 2021, Microsoft offered federal agencies “free” security upgrades, along with consultants to install them. According to former employees interviewed by ProPublica, the company’s strategy was simple. Once the upgrades were installed, agencies would find it too expensive and disruptive to switch, leaving them little choice but to pay once the free period ended. When agreements came up for renewal in 2022 and 2023, every Defense Department customer accepted the offer, including every military branch.⁴

This is a strategy motivated by a profit imperative; it is not sensitive to warnings about the risks AI can introduce when integrated into government systems. Last month, the New York Public Service Commission launched an inquiry into public utility use of AI, requiring electric, gas, and water utilities to disclose how they use AI, warning that the technology is vulnerable to “hallucinations, algorithmic biases, transparency issues, data privacy concerns,

¹ Goldman Sachs, “Global AI Investment Is Forecast to Exceed \$1 Trillion in 2026,” August 7, 2026, <https://www.goldmansachs.com/insights/articles/global-investment-is-forecast-to-exceed-1-trillion-in-2026>.

² David Crawford et al., “New Innovation Is Required to Fund AI’s \$6 Trillion Buildout,” in *Technology Report 2026* (Bain & Company, September 29, 2026), <https://www.bain.com/insights/new-innovation-is-required-to-fund-ais-6-trillion-buildout-technology-report-2026/>.

³ U.S. General Services Administration, “Buy AI,” last modified September 10, 2026, <https://www.gsa.gov/artificial-intelligence/buy-ai>.

⁴ Renee Dudley, “Microsoft’s “Free” Plan to Upgrade Government Cybersecurity Was Designed to Box Out Competitors and Drive Profits, Insiders Say,” *ProPublica*, November 15, 2024, <https://www.propublica.org/article/microsoft-white-house-offer-cybersecurity-biden-nadella>.

misconfiguration errors, cybersecurity attacks, and functional brittleness,” with potential consequences for the safety and reliability of the state’s critical infrastructure. Regulators conceded that the extent and scope of AI systems in New York utility operations remains unclear.⁵ In other words, AI is already embedded in the systems that keep the lights on and the water running in New York City, and the public does not know where or how.

These examples reveal a pattern. AI companies are embedding their products in public institutions faster than agencies, let alone the public, can understand what is being installed.⁶ Once in place, these systems are difficult to audit, costly to remove, and controlled by companies whose obligations run to their shareholders, not to New Yorkers.⁷ New York City must prioritize New Yorkers when deciding how AI enters the services its residents depend on.

To start, New York City must **treat every point of integration as an affirmative choice, and every choice as leverage**. New Yorkers should not be the testing ground for an industry’s experiments. This approach can and should be extended to every city agency.

In practice, this means no AI system should enter City operations without an affirmative decision by the City. Before procurement, an agency should be required to demonstrate that the tool is fit for purpose, works as intended, and is a valid way to solve the problem in front of it. The agency should also assess what the tool could cost New Yorkers, and that assessment must cover city workers and the communities most targeted by surveillance and policing. Contracts should limit data collection to what is reasonably necessary and bar sharing it across agencies. Any new use of a tool, whether in a new department, on a new population, or with a new kind of data, should require new approval.⁸

The City must also apply increased scrutiny to promotional offerings from AI companies. Agencies should review free and heavily discounted offerings, like the \$1 deals extended to federal agencies, with the same rigor as paid contracts. While these deals may appear low-cost on the surface, their cost for the City—in potential compliance errors and risks to the public—is deferred rather than waived. This level of rigor must also be applied to AI that is introduced into city systems after a contract has been signed. AI is increasingly deployed through the back door. Existing vendors are adding AI features to software after customers have already integrated into its operations, often through routine updates or contract renewals, without any affirmative

⁵ Robert Walton, “New York Audits Utility AI Use, Cites Risk in ‘Growing Dependency,’” *Utility Dive*, September 24, 2026, <https://www.utilitydive.com/news/new-york-audits-utility-ai-use-cites-risk-in-growing-dependency/831242/>.

⁶ Taraaz and Collaborative Research Center for Resilience, *Key Considerations When Procuring AI in the Public Sector* (August 2025), https://static1.squarespace.com/static/5d159d288addab0001036c45/t/6890f9066bf93951bedd9485/1754331401682/AI_Procurement_Taraaz_CRCR_2025.pdf; See also U.S. General Services Administration, “Buy AI.”

⁷ See Taraaz and Collaborative Research Center for Resilience, *Key Considerations When Procuring AI in the Public Sector*; and Dudley, “Microsoft’s “Free” Plan.”

⁸ See Taraaz and Collaborative Research Center for Resilience, *Key Considerations When Procuring AI in the Public Sector*.

decision by the agency using it.⁹ A contract renewal is a point of integration, and City agencies and departments should treat it as a procurement decision subject to the same scrutiny as a new purchase.

2. The AI Harms New Yorkers Already Face

While many of the risks of algorithmic systems in government are still taking shape as adoption grows, the private sector, where AI has spread with far less oversight, shows us what comes next, and how New Yorkers are already paying the price.

The same companies courting city agencies have also been remarkably successful at defining what counts as an AI safety problem. By focusing public debate on speculative, long-term catastrophes, this framing diverts attention from the tangible harms algorithmic systems are already causing for New Yorkers.

Algorithmic systems shape and make decisions about New Yorkers' lives every day. When New Yorkers walk into Madison Square Garden, the "World's Most Famous Arena," AI-enabled facial recognition technologies scan their faces and check them against the venue's internal watchlists. The facial recognition system has been used to enforce bans against hundreds of people, including lawyers whose firms were in litigation with MSG, many of whom had no role in the cases themselves. It has been used to track a trans woman's movements at the Garden, including her trips to the bathroom, for two consecutive years.¹⁰ MSG has even compiled a dossier on prominent activists campaigning against facial recognition technology and biometric surveillance.¹¹

These technologies watch New Yorkers more closely, more accurately, and across more connected systems than ever before, unbeknownst to most New Yorkers. When New Yorkers

⁹ Kareem Anderson, "Microsoft Says It'll Force Install Microsoft 365 Copilot on Windows 11 with MS 365 Business in the Next 30 Days," *Windows Latest*, June 22, 2026, <https://www.windowslatest.com/2026/06/22/microsoft-says-itll-force-install-microsoft-365-copilot-on-windows-11-with-ms-365-business-in-the-next-30-days/>; See also Google, "The Best of Google AI, Now Included in Workspace Business and Enterprise Plans," *Google Workspace Updates* (blog), January 15, 2025, <https://workspaceupdates.googleblog.com/2025/01/expanding-google-ai-to-more-of-google-workspace.html>.

¹⁰ Noah Shachtman and Robert Silverman, "The Shocking Secrets of Madison Square Garden's Surveillance Machine," *Wired*, April 17, 2026, <https://www.wired.com/story/madison-square-garden-jim-dolan-surveillance-machine/>.

¹¹ Joseph Cox, "Madison Square Garden Made Dossier on Activists Who Opposed Facial Recognition," *404 Media*, June 23, 2026, <https://www.404media.co/madison-square-garden-made-dossier-on-activists-who-opposed-facial-recognition/>.

apply for a job,¹² seek housing assistance,¹³ seek care in an emergency room,¹⁴ go to a concert, log onto a delivery app,¹⁵ or walk through a subway turnstile,¹⁶ their presence is perceived, analyzed, and recorded. This surveillance presents glaring safety risks, yet does not enter industry-led discussions around AI safety and risk.

New Yorkers do not need to imagine how automated systems can harm them. More than one in five New Yorkers work gig jobs.¹⁷ As the city has found, gig platforms have used their control over their algorithmic systems to undercut protections that workers have fought for. When New York City's landmark minimum pay standard for app-based delivery workers took effect, in the words of Mayor Mamdani, the delivery apps "buried the tip button to pad their own bottom line," costing delivery workers an estimated \$550 million in tips.¹⁸

It's important to remember that these are not simply technical problems that will go away as the technology improves. The harmful effects of the unbridled integration of AI-powered tools must be evaluated within the social and economic context in which they will be deployed. When these technologies are deployed by companies prioritizing their interests over our communities, the exploitation will only intensify. In many cases, that means exacerbating existing inequities in public safety. New Yorkers living in areas at greater risk of stop-and-frisk by police, areas with a higher proportion of non-white residents, are also more exposed to invasive facial recognition

¹² Office of the New York State Comptroller, "DiNapoli: New Yorkers Deserve a Transparent Hiring Process When Artificial Intelligence Is Used To Vet Their Job Applications," press release, December 2, 2025, <https://www.osc.ny.gov/press/releases/2025/12/dinapoli-new-yorkers-deserve-transparent-hiring-process-when-artificial-intelligence-used-vet-their>.

¹³ New York City Office of Technology and Innovation, *Agency Compliance Reporting of Algorithmic Tools: Calendar Year 2025* (March 2026), <https://www.nyc.gov/assets/oti/downloads/pdf/reports/LL35%20Report%202025%20-%20Final%20-%202026-03-27.pdf>.

¹⁴ BronxCare Health System, "Emergency Medicine," accessed October 1, 2026, <https://www.bronxcare.org/our-services/emergency-medicine>; See also Northwell Health, "Northwell Health to Deploy Ambient AI Technology to Ease Clinicians' Workload," press release, <https://www.northwell.edu/news/the-latest/ambient-ai-tech-bridge-to-ease-clinician-note-taking>.

¹⁵ The City University of New York and Worker's Justice Project, *Driven by Algorithms: Unsafe Work in NYC Food Delivery*, data brief (September 2026), <https://cunyurbanfoodpolicy.org/resources/report/driven-by-algorithms/>; Sam Levine, "Reclaiming Worker Power in the Age of AI," *The Nation*, October 2, 2026, <https://www.thenation.com/article/politics/reclaiming-worker-power-in-the-age-of-ai/>; Derek Kravitz, "Instacart's AI-Enabled Pricing Experiments May Be Inflating Your Grocery Bill, CR and Groundwork Collaborative Investigation Finds," *Consumer Reports*, December 9, 2025, <https://www.consumerreports.org/money/questionable-business-practices/instacart-ai-pricing-experiment-inflating-grocery-bills-a1142182490/>.

¹⁶ Kevin Collier, "NYC subway using AI to track fare evasion," *NBC News*, July 20, 2023, <https://www.nbcnews.com/tech/tech-news/nyc-subway-using-ai-track-fare-evasion-rcna93045>; Sophia Fox-Sowell "AI Cameras Will Flag 'Problematic Behavior' in New York Subways, MTA Says," *StateScoop*, April 30, 2025, <https://statescoop.com/mta-ai-cameras-ny-subway/>.

¹⁷ Rachel Swaner, *Surveilled and Directed: Algorithmic Management and Control in New York's Gig Economy* (Community Service Society of New York, April 2026), https://smhttp-ssl-58547.nexcesscdn.net/nycss/images/uploads/pubs/040726_Annual_Survey_Gig_Economy_V5.pdf.

¹⁸ New York City Office of the Mayor, "Mayor Mamdani Delivers More Than \$100 Million for Delivery Workers Since January," press release, July 29, 2026, <https://www.nyc.gov/mayors-office/news/2026/07/mayor-mamdani-delivers-more-than--100-million-for-delivery-workers>.

technology.¹⁹ Making that system more accurate doesn't make those New Yorkers safer. It makes the targeting more precise.

The city must reassert its jurisdiction in response to the rapid rollout of AI tools in the private sector. Companies that break the law and harm New Yorkers should face consequences, and the city already has the tools to impose them. The City's own findings on how delivery platforms obscured tip options on their apps after the minimum pay standard took effect show that municipal enforcement can identify algorithmic evasion when it looks for it.

Enforcement after the fact, however, cannot address harms that New Yorkers are unaware of in the first place. **The City should adopt bright-line rules that prohibit the most invasive uses of these technologies outright.** This includes passing ongoing Council efforts to prohibit facial recognition in places of public accommodation, and to limit its use in residential buildings. It should also include a ban on facial recognition and biometric surveillance in policing, with particular attention to the loopholes and workarounds through which such bans have been evaded elsewhere.

Where algorithmic systems are used to make or inform decisions about New Yorkers, the City must ensure that people know and have a meaningful choice. New Yorkers should be notified when AI is used in a decision that affects them. In the highest-stakes contexts, the question is not how to mitigate these systems but whether they should be used at all. These protections must extend to New Yorkers as workers and consumers. The City should ban the use of surveillance pricing and wage-setting, practices in which companies use personal data to set individualized prices and pay.

3. The Concentration of Power in AI—and the Single Points of Failure it Engenders—is a Safety Risk for New York City

There is no AI without Big Tech, without the handful of firms whose infrastructure sits beneath nearly everything we use.²⁰ AI developers depend on the same compute, cloud, and distribution companies that New York City services, hospitals, and small businesses run on.

As AI makes cyberattacks faster and harder to detect,²¹ these attacks are increasingly dangerous because of how concentrated the industry has become. The International Monetary Fund has

¹⁹ Amnesty International, "USA: Facial Recognition Technology Reinforcing Racist Stop-and-Frisk Policing in New York—New Research," February 15, 2022, <https://www.amnesty.org/en/latest/news/2022/02/usa-facial-recognition-technology-reinforcing-racist-stop-and-frisk-policing-in-new-york-new-research/>.

²⁰ Kate Brennan, Amba Kak, and Sarah Myers West, *Artificial Power: 2025 Landscape Report* (AI Now Institute, June 3, 2025), <https://ainowinstitute.org/publications/research/ai-now-2025-landscape-report>.

²¹ Lizzie Short, "AI and the Future of Cybersecurity," Harvard Extension School (blog), July 21, 2025, <https://extension.harvard.edu/blog/ai-and-the-future-of-cybersecurity/>.

warned that a single flaw can cascade across energy, telecommunications, and public services.²² The more deeply this technology is entrenched in our systems, the further that damage can spread.

In August 2025, Anthropic reported that a lone hacker used Claude Code to run nearly every stage of an extortion campaign against at least 17 organizations in a single month, among them healthcare providers, emergency services, and government agencies, exposing medical records, financial data, and government credentials.²³ These are the same kinds of institutions New Yorkers depend on every day, and every dollar they spend defending against AI-enabled attacks is a dollar diverted from care, services, and staff.

Even routine failures beyond cyberattacks can trigger global outages across thousands of services. In the span of a single month last fall, AWS, Cloudflare, and Azure all experienced internal errors that impacted the services New Yorkers depend on every day, from Slack at work,²⁴ to Uber and NJ transit's digital services for getting around,²⁵ to Costco for household essentials.²⁶ The industry that is raising hundreds of billions of dollars to build, deploy, and integrate AI is the actor best positioned to bear the costs of making it safe. Instead, those costs are being absorbed by the people and institutions with the fewest resources to absorb them.

New York City cannot, on its own, resolve the concentration of power in the technology industry. It is, however, an enormous purchaser of technology, spending billions of dollars through just a handful of citywide IT contracts,²⁷ and its procurement decisions either deepen that concentration or resist it. The City should begin by asking a question it has too often skipped: whether it wants its public infrastructure built on the infrastructure of a handful of private companies.

Finally, the costs of failure should be borne by those best positioned to prevent it. City contracts should **require vendors to meet rigorous security standards, disclose breaches and outages promptly, and assume financial liability when their systems fail.**

²² Tobias Adrian, Tamas Gaidosch, and Rangachary Ravikumar, "Financial Stability Risks Mount as Artificial Intelligence Fuels Cyberattacks," *IMF Blog*, May 7, 2026, <https://www.imf.org/en/blogs/articles/2026/05/07/financial-stability-risks-mount-as-artificial-intelligence-fuels-cyberattacks>.

²³ Anthropic, *Threat Intelligence Report: August 2025* (August 2025), <https://www-cdn.anthropic.com/b2a76c6f6992465c09a6f2fce282f6c0cea8c200.pdf>.

²⁴ Leo Sands and Naomi Nix, "Amazon Cloud Outage Disrupts Service of Major Apps," *The Washington Post*, October 20, 2025, <https://www.washingtonpost.com/technology/2025/10/20/aws-outage-amazon-fortnite-snapchat-offline/>.

²⁵ Michelle Chapman, "Cloudflare Outage Impacts Thousands, Disrupts Transit Systems, ChatGPT, X and More," ABC7 New York, November 18, 2025, <https://abc7ny.com/post/cloudflare-outage-disrupts-chatgpt-other-internet-services/18169305/>; James Peckham, "Cloudflare Outage Takes Down ChatGPT, Uber, X, More: Here's What Happened," *PCMag*, November 18, 2025, <https://www.pcmag.com/news/cloudflare-outage-takes-down-chatgpt-uber-x-more-heres-what-happened>.

²⁶ Associated Press, "Microsoft Deploys a Fix to Azure Cloud Service That's Hit with Outage," *AP News*, October 29, 2025, <https://apnews.com/article/microsoft-azure-down-detector-service-disruption-0deffbd09c09ca4640c2f5452a9e483e>.

²⁷ Office of the New York City Comptroller, *The Monty Hall Contracts: Unchecked Spending Across the City's Master Agreements* (May 6, 2026), <https://comptroller.nyc.gov/reports/the-monty-hall-contracts-unchecked-spending-across-the-citys-master-agreements/>.

Conclusion

New York City still holds the leverage to decide which systems enter its agencies, to set bright-line rules for how AI may be used on the people who live and work here, and to make the industry bear the costs of its failures. I urge the Council to use that leverage now.