

How New York City Can Use Artificial Intelligence to Improve Quality of Life

Josh Appel

Policy Analyst
Manhattan Institute

Introduction

In May 2026, New York City Mayor Zohran Mamdani created a Commission on Government Efficiency (COGE) “to modernize City government, improve service delivery, and strengthen accountability to New Yorkers.”¹ Improved service delivery is a timely goal. In New York City today, violent crime rates remain relatively low,² but visible signs of disorder—graffiti, public indecency and urination, homelessness in the transit system, and trash buildup—are increasing. Since Mamdani took office, the number of quality-of-life 911 calls has risen by 13%, with transit-related calls increasing by 36%. Public urination, graffiti, and disorderly conduct are all up significantly as well.³ A mere 34% of New Yorkers rated the city’s quality of life as “excellent or good” in 2025, a decrease from 51% in 2017.⁴ If that continues, New Yorkers will feel increasingly unsafe on the streets and subways.

About Us

The Manhattan Institute is a community of scholars, journalists, activists, and civic leaders committed to advancing economic opportunity, individual liberty, and the rule of law in America and its great cities.

If Mayor Mamdani truly wants to improve service delivery to improve the quality of life for New Yorkers, the city government should embrace the power of artificial intelligence (AI). Decades ago, New York City showed how data-driven governance can radically transform urban life. From the NYPD’s CompStat system in the 1990s to the FDNY’s FireCast predictive inspection model and Mayor Bloomberg’s creation of 311 and the Mayor’s Office of Data Analytics (MODA), the city pioneered the use of large-scale data collection to improve public safety, allocate resources more efficiently, and prevent problems before they escalated.

AI presents an opportunity to build on that legacy. Unlike earlier systems, which were largely limited to single categories of data, modern AI can process massive quantities of text, video, audio, and real-time information simultaneously, allowing the city to identify patterns across quality-of-life indicators with unprecedented speed and sophistication. By leveraging the vast data that the city *already* collects through NYC Open Data and existing



city agencies, New York can create predictive models focused on issues ranging from reckless driving and illegal vending to nightlife disorder, sanitation overflow, subway congestion, and homeless encampments. Just as CompStat and FireCast revolutionized policing and fire prevention, AI-powered analytics could usher in a new era of proactive governance in which city services are deployed more intelligently, efficiently, and preventively to improve public order and neighborhood quality of life.

How Data-Driven Policies Changed New York Once Before

New York City's embrace of data-driven governance emerged from a series of institutional experiments over two decades. The story begins in the early 1990s during a seemingly intractable crime crisis and concludes with Mayor Bloomberg's ambitious effort to build a citywide analytics infrastructure that would outlast his administration. Over that period, NYC built a model for data-driven governance, built on data collection and predictive allocation. AI models can build on that foundation.

In 1994, Police Commissioner William J. Bratton launched CompStat, a system that required precinct commanders to map and track crime incidents in real time and then defend their response strategies at weekly accountability meetings. The results were striking. Between 1993 and 2000, murders in NYC fell by nearly 70%, from 1,927 per year to 673.⁵ Data-driven policing also drove down property crime and other violent crime.⁶ CompStat was replicated in cities across the country, becoming one of the most influential public-management innovations of the late twentieth century.

Criminologists continue to debate how much credit belongs to CompStat, versus other factors.⁷ Critics of the program have also pointed to investigations and whistleblower accounts revealing credible evidence of statistical manipulation in some precincts, suggesting that some of the reported declines in crime may have been overstated. These criticisms are legitimate, and they offer lessons that will be relevant as any future data initiative is implemented. Nevertheless, nearly all analysts agree that CompStat was a key part of data-driven policing, which was likely responsible for the dramatic fall in the murder rate.⁸

FDNY also used data to help with the proper division of scant resources. In 2007, after an immense fire killed two firemen despite a quick response time, the department wanted to improve the building inspection process to prevent sizable fires in the first place. Rich Tobin, assistant chief of the Bureau of Fire Prevention, worked with IBM to build the Risk Based Inspection System, which could determine a building's level of fire risk. This allowed inspectors, who were in short supply, to devote their time to the buildings that most needed inspection. The main component of the system was FireCast, which assigned risk factors based on data from various city agencies, including building age, vacancy status, prior violation records, nearby utility complaints, and socioeconomic indicators from census data—inputs that no single agency had ever synthesized. The system flagged roughly 15% of the city's buildings as high-risk and directed inspectors toward them accordingly. When FireCast 2.0 launched in 2013, FDNY officials reported that 16% of fires were in buildings that had been inspected over the previous 90 days.⁹ It was a proof of concept for something the city had never quite managed: using algorithmic cross-agency pattern recognition not just to respond to problems, but to anticipate them.

Alongside FireCast, Mayor Bloomberg’s team worked on its own data-collection and sharing plan. In 2003, Bloomberg launched 311, creating an influx of millions of data points. Bloomberg appointed Mike Flowers, chief analytics officer for Enigma, to run New York’s first data analytics office. Together with the mayor, Flowers worked to prevent and respond to public-safety issues. One early success involved illegal basement conversions. By combining 311 complaints, Department of Buildings permit data, utility consumption records, and census information, Flowers’s team was able to identify buildings with high probabilities of illegal occupancy—a major fire risk—long before inspectors would have found them through routine checks. The Mayor’s Office of Data Analytics, formally established in 2010 and codified in the city charter in 2018, institutionalized this approach. MODA’s mandate was explicitly cross-agency: to break down data silos and allow city government to see patterns that no individual department could detect on its own.

Collecting more data also allows for faster and easier pattern recognition. Better pattern recognition allows resources to be better allocated at the outset to prevent big problems down the line, thus reducing the expenditure required to address them. FireCast, for example, reduced the number of massive fires—which would require hundreds of firemen for many hours at increased risk—because the algorithm could use large amounts of data to flag buildings at greatest risk of fire beforehand. Data alone cannot fix problems but do allow decision-makers to prioritize resources more effectively.

My colleague Nicole Gelinas wrote extensively of the “Fourth Urban Revolution,”¹⁰ charting the success of data-driven resource allocation and policy. Writing a decade ago, Gelinas argued that the city’s data infrastructure, while impressive, remained fragmented and in need of an update. AI large language models offer an opportunity to do just that. Where CompStat could measure crime and FireCast could assess building risk, LLMs can synthesize text, images, audio, and real-time sensor data across every agency simultaneously—and do so at a speed and scale that earlier systems could not approach. The foundation built in the Bloomberg era is still in place. The question is whether we will use new tools to build on it.

Artificial Intelligence and Data Collection Today

As the name implies, large language models such as ChatGPT or Claude operate best with robust data sets. Luckily, New York City already collects a vast amount of data. Many New Yorkers are unaware that the city currently collects millions of data points and uploads them for public use through NYC Open Data. Every city-controlled input is meticulously recorded. Turnstile usage, traffic stops and violations, tree data, squirrel censuses, permit use, garbage collection and overflow, every 311 or community-board complaint, prison inmate assaults and deaths—the list is almost endless. All these data are open and transparent through NYC Open Data, part of the Office of Technology & Innovation (OTI).¹¹

As of now, Open Data is regularly used by nongovernmental organizations. For example, Stop the Chop, an independent watchdog group, used Open Data to raise awareness about helicopter noise.¹²

NYC Open Data was signed into law in 2012 by Mayor Bloomberg, thanks to the efforts of then-city councilmember Gale Brewer. Local Law 11 of 2012 states that “it is in the best interest of New York City that its agencies and departments make their data available online using open standards. Making city data available online using open standards will make the operation of city government more transparent, effective and accountable to the public. It will streamline intra-governmental and inter-governmental communication and interoperability, permit the public to assist in identifying efficient solutions for government, promote innovative strategies for social progress, and create economic opportunities.”¹³ The Open Data Law expanded upon the 1974 Freedom of Information Law, which required agencies to provide data upon request. In 2015, the de Blasio administration sought to make government data even more accessible by increasing the number of available datasets beyond the original 1,300 and by improving the platform’s usability.¹⁴ One major result was the release of all city data by 2018. Today, Open Data has almost 2,412 total datasets with 6,398,693,616 rows of data—and counting.¹⁵

Open Data, though, is not without its problems. For one, many datasets are outdated because of a lack of automation, meaning that data are not updated in real time, leading to collection backlogs. In 2024, Rachael Fauss, senior policy advisor for Reinvent Albany, testified that it took OTI 12 years to automate 435 datasets while 437 remain unautomated.¹⁶ A common approach to automating large datasets is an Extract, Transform, Load (ETL) pipeline. For example, with an ETL pipeline, every speed-camera violation would automatically appear on Open Data within minutes (because collection is already a digital input). Without an ETL pipeline, in order for something to appear on Open Data, employees need to contact data officers at each department and then work through and format the data themselves. Currently, many datasets do not use an ETL pipeline, and data collection from agencies involves an unnecessary and cumbersome two-step process. In 2017, Local Law 251 required each city agency to report whether their data *could* be automated but did not mandate automation. Today, many datasets remain unautomated.¹⁷

Another major problem with Open Data is ease of use, or lack thereof. Data do not live on a visual dashboard but on a spreadsheet with thousands of rows and columns. This means that the average person has no way of using any of these data to assess whether city services are being delivered adequately.

In 2025, David Tussey, former executive director of the Department of Information Technology and Telecommunications, together with Jun Yan, professor in the Department of Statistics at the University of Connecticut, illustrated the problems with Open Data with a case study of 311 data.¹⁸ Even to assess the data, Tussey and Yan had to download over 3 gigabytes of information with millions of rows. They found that there was no uniform or consistent data style and no internal mechanism to evaluate and correct errors. The spreadsheet contained incorrect formatting, fake zip codes, and null information, and time stamps did not account for daylight saving time, possibly skewing results. To improve the usability of its data, Tussey and Yan recommended that the city consult with experts who could “help design and refine curation processes, such as creating algorithms for automatic data cleaning, identifying inconsistent data entries, and ensuring the dataset(s) structure is maintained over time.”

AI models excel in automating and streamlining data, correcting mistakes, eliminating inconsistent formatting, and producing data visualizations. AI is thus a valuable resource to improve the data-collection process, which is a first step toward improved pattern recognition and resource allocation.

How AI Data Collection Can Improve the Mayor's Management Report

Another area where AI can be a valuable tool is the Mayor's Management Report (MMR), released by the mayor every fiscal year. It reports on how well the city is meeting performance benchmarks for service delivery. However, as is often true with government services, the focus of the report is sometimes on how much the city is spending or doing, rather than on the outcomes the city has achieved.¹⁹

The section on FDNY's performance in the 2026 report, for example, begins by touting not the number of fires reduced but the increased number of inspections: "In the first four months of Fiscal 2026, fire companies conducted over 13 percent more mandatory inspections by uniformed personnel (from 9,325 to 10,580) than over the same period in Fiscal 2025."²⁰ There is no analysis of the budget-service ratio, whether money is well spent, how well the FDNY did in putting out fires, or New Yorkers' satisfaction with their performance. The report alludes to several response-time data metrics but only indicates whether performance is better or worse than last year for any given metric, rather than providing the raw data. In addition, because MMR is published as a PDF, it is impossible to analyze the data by, for example, filtering different zip codes, boroughs, or years.

MMR, like Open Data, would benefit from the integration of LLMs. Rather than publishing metrics on a spreadsheet, MMR can use an LLM combined with an ETL pipeline to get updated data, a consistent format, and the ability to visualize that data by year, borough, or zip code. Instead of spending hours collecting inputs, the *outcomes* can be automatically registered each day and kept up to date. LLMs also allow for better categorization of data. For example, rather than simply measuring how many community-board complaints or 311 calls there were in a given year, we can filter complaints and calls for different issues and see whether departments are helping the community or not. These data can, in turn, be easily visualized with built-in filters, graphs, and maps.

How AI Can Improve Pattern Recognition

MODA—which uses analytics “to aggregate and analyze data from across City agencies, to more effectively address crime, public safety, and quality of life issues”²¹—is in need of an update.

With LLMs, it is easier than ever to use data to recognize patterns. To understand why, it helps to understand what distinguishes LLMs from earlier data tools used by the city. Systems like CompStat or FireCast were purpose-built: they ingested a defined set of structured inputs—incident reports, building codes, inspection records—and produced outputs within a narrow domain. They were powerful within their lane but blind beyond



it. LLMs are structurally different. They are general-purpose reasoning engines trained across virtually every domain of human knowledge, capable of understanding context, drawing analogies, and synthesizing information across subject areas that have no obvious connection. For city government, that distinction matters enormously.

The first major advantage over previous systems is multimodal processing. Traditional data analytics pipelines require clean, structured inputs such as rows and columns and standardized fields with consistent formatting. These systems cannot integrate unstructured data, such as a 311 complaint written in plain English, a photograph of a cracked sidewalk, a noise complaint recorded as audio, or a surveillance feed flagging unusual crowd behavior. LLMs were built for exactly that kind of messy, real-world information. A modern multimodal LLM can read a building inspector's handwritten notes, cross-reference them against a violation database, examine a photo of the premises, and synthesize all three into a single risk assessment. That would take hours of work for any human analyst, which is why it is not something we have ever done. The city already generates millions of unstructured data points every day across hundreds of agencies. With AI, it can actually use those data in meaningful ways.

The second advantage of LLMs is speed at scale. The bottleneck in data-driven governance is not the availability of data but the capacity to synthesize it at *both* speed and scale to create a positive predictive model. The entire advantage is something that a human simply cannot do: inspect millions of data points to construct predictive equations with ease.

Taken together, multimodality and speed create a compounding effect. The value of any single data source is multiplied when it can be combined with others, and LLMs make that combination not just possible but routine. Whereas CompStat could measure crime patterns within NYPD's data silo, and FireCast could assess building risk within FDNY's, an LLM-powered system can span both, and every other city agency simultaneously, finding correlations that no single-domain system would ever detect. A spike in 311 noise complaints in a given neighborhood, combined with increased foot traffic detected via transit data, combined with a cluster of recent parole completions flagged in DOC records, might together predict an uptick in violent incidents days before it materializes. No CompStat precinct commander would have had the bandwidth to synthesize those three streams, but an LLM could do it automatically.

Consider how a next-generation FireCast built on an LLM foundation might actually look. The current system scores buildings on roughly 7,500 risk factors drawn from 17 different agencies, already a marvel in technological innovation. But an AI-powered successor could ingest more varied types of real-time data from even more agencies—not just building violations and inspection records but also live weather data (high winds and low humidity dramatically increase fire spread), utility outage reports that may indicate electrical faults, social-media posts flagging visible smoke or smells, geospatial data on hydrant access, and street-width data that could affect engine response time and staffing levels at nearby firehouses on any given shift. The system could update risk scores in real time rather than semiregularly, flag buildings above a certain risk threshold, and generate plain-language briefings about the highest-risk exposures for battalion chiefs each morning. That is not a vision of some distant AI future because these components exist today. What is missing is the institutional will and interagency coordination to assemble them.

The uses of LLM-powered data analysis extend well beyond fire prevention. The same model could be applied to any area in which data are currently fragmented and the city's response is reactive, rather than predictive: pothole repair dispatched after a complaint rather than predicted from freeze-thaw cycle data and road-age models; homelessness

outreach triggered by visible encampments rather than by early indicators like eviction filings, utility shutoffs, and emergency-room visits; sanitation routes operated on fixed schedules rather than real-time fill-level sensors and complaint density. In each case, most of the relevant data already exist. LLMs provide only the connective tissue to turn those data into foresight rather than hindsight. Being able to detect patterns, update in real time, and process fundamentally different types of information simultaneously means better resource allocation across every agency and, ultimately, a safer and more functional city.

Examples

There are too many possible applications of LLM-powered data analysis to cover in this issue brief, but a few examples will help flesh out the concept and the type of approach necessary to create quality-of-life improvements.

1. **Reckless driving:** Drag racing and reckless night driving (doughnuts in an intersection, for example) are dangerous, antisocial behaviors that corrode neighborhood tranquility, especially for families with young children. At present, the city learns about a drag race from 311 complaints filed after it has already happened. The inputs needed to anticipate one, however, already exist: 311 complaint texts, speed-camera placement and violation records, traffic volume and street geometry, NYPD incident narratives and patrol schedules. What has been missing is the capacity to convert those inputs into usable form. Complaints arrive as free text, camera output as video, noise reports as audio, and police accounts as narrative prose. None of it arrives in the rows and columns that a statistical model requires, which is why this analysis has never been attempted at scale. Multimodal AI closes that gap. A vision model can process camera footage to count vehicles, flag repeat plates, or detect the signature of a burnout at an intersection. An AI can read tens of thousands of 311 complaints and incident reports and tag each by behavior, time, and location. Together, they turn unstructured municipal reporting into structured variables at a cost per record that finally makes citywide analysis feasible. Once those variables exist, a conventional predictive model does the forecasting: in a corridor with X uninterrupted blocks between traffic controls, Y speed cameras, and Z officers assigned to a weekend overnight tour, the likelihood of a racing event rises by a measurable and testable amount. The innovation here is not the statistical method of regression analysis, which was available to Mike Flowers in 2011; the innovation is the volume and type of data and how efficiently the data are cataloged. An AI then makes the result usable to the people who must act on it, generating a plain-language weekly briefing for a precinct commander showing which corridors have crossed a risk threshold, and answering follow-up questions without anyone writing a database query. The output is operational: where to assign patrol on a Saturday night; and where an additional speed camera would be best.
2. **Trash collection:** The city already collects all permit information across each borough. The city also has adequate data to know when there is trash overflow (NYC Department of Sanitation [DSNY], 311 call patterns). We can create a predictive model that shows the pattern of permit type and trash overflow. We would then be able to say that Permit X likely leads to Y chance of trash overflow, thus allowing DSNY to take proper action (sending an extra pickup) if deemed necessary.



3. **Subway scheduling:** Using AI, the Metropolitan Transit Authority can take existing data to figure out which subways are most and least congested and realign schedules accordingly. This could also help reduce subway traffic.
4. **Nightlife-related disorder:** The city can create predictive models of where nightlife-related disorder is most likely to emerge before conditions deteriorate, using data on liquor licenses, 311 noise complaints, sanitation overflow reports, NYPD incident data, EMS calls, pedestrian traffic patterns, taxi and rideshare congestion, and late-night subway usage. These datasets make it possible to determine which combinations of nightlife density, operating hours, street design, police visibility, and sanitation schedules correlate most strongly with spikes in disorder, noise, congestion, and public-safety incidents. For example, the system may determine that corridors containing a high concentration of bars open past midnight, narrow sidewalks, low late-night transit accessibility, and insufficient sanitation pickup schedules are substantially more likely to experience increases in noise complaints, litter accumulation, assaults, and public intoxication. The city could then proactively allocate resources before conditions worsen by increasing patrol presence, adjusting sanitation schedules, redesigning sidewalks, staggering closing hours, or limiting new licenses within already-saturated/high-risk areas.
5. **Illegal vending:** Right now, the city typically treats illegal vending as a matter for isolated enforcement, despite its broader impact on public order and neighborhood functionality. Illegal vending frequently correlates with sidewalk congestion, sanitation accumulation, blocked subway entrances, unsafe pedestrian flow, counterfeit-goods sales, and repeated 311 complaints. The city already collects substantial data on these conditions, including pedestrian traffic flows, MTA ridership, sanitation records, prior summons locations, geospatial sidewalk measurements, event schedules, and complaint histories. Using AI-powered predictive analysis, the city could identify where vending-related congestion and disorder are most likely to emerge and determine which interventions are most effective at preventing deterioration. AI systems could recognize recurring spatial and temporal patterns showing that certain transit hubs, commercial corridors, or event-adjacent areas become especially vulnerable to overcrowding, trash overflow, or blocked pedestrian movement under specific conditions. These models would also help show where illegal, counterfeit, or resold licenses are mostly likely to show up. The city could then proactively deploy the proper sanitation resources, establish temporary legal vending zones, remove illegal vendors, and increase targeted enforcement.
6. **Quality of life:** We can use all possible data to create quality-of-life index scores for every district, thus magnifying which areas are most underserved and in need of repair.
7. **311:** We can combine AI data analysis with documented human response to figure out which 311 calls should be taken most seriously and which, such as calls that are likely fraudulent, should be de-prioritized.
8. **Homelessness:** Another example is a homeless encampment prediction model. Right now, city agencies tend to respond episodically after encampments become entrenched or dangerous. But recurring patterns certainly exist, involving transit hubs, vacant storefronts, robust scaffolding, lighting conditions, shelter proximity, weather changes, and prior enforcement history. An AI model could identify which public spaces are most vulnerable to becoming chronic disorder zones and recommend preventive interventions such as increased outreach, law enforcement, sanitation deployment, or redesign of physical space.

Interagency Data Sharing

Although interagency sharing of data can raise legal questions about data privacy, those concerns are not implicated in these proposals. The data necessary for most of the applications described in this brief are already public and digitized. Any person can open Claude on a browser today and ask for a data visualization of trash pickup or subway schedules, combining those prompts with all sorts of other creative inputs. If any applications do raise substantial legal issues, it is possible to form interagency data-sharing agreements, such as the 2008 Inter-Agency Data Exchange Agreement, which created a memorandum of understanding allowing HHS to share data with certain agencies such as DHS or NYCHA to improve services.²² Theoretically, other agencies interested in this technology could sign on to a similar agreement to be applied for broader use.

What Kind of Model to Build

To get the most from AI-powered data tools, the mayor's office should use purpose-built internal infrastructure rather than rely on consumer-facing commercial tools. The most practical path is not to build a model from scratch but to deploy a self-hosted or enterprise-licensed version of an existing foundation model, fine-tuned on city-specific data and governed by city-specific access controls. JPMorgan, for example, built an internal model called LLM Suite (later expanded into a broader AI platform) specifically calibrated to its operational needs and customer data.²³ The city should take a similar approach: take a capable foundation model, restrict it to approved data sources, train it on municipal workflows, and run it on New York's infrastructure. This architecture has several advantages. First, it keeps sensitive government operations off public-facing servers. In addition, it allows the model to be customized for specific agency needs. A fire-risk scoring model has different requirements from a pothole-prediction model, for example, and both would be different from taking an existing AI from the open market, even one tailored to generic cities. Most important, such a model gives the government full control—allowing for immediate shutdown or rollback if the system produces harmful outputs—without the need for a third party.

Privacy Concerns

Any proposal to expand data-driven governance should take seriously the civil liberties concerns that have been raised in connection with certain predictive policing initiatives. The most divisive predictive policing programs, however, are those focused on identifying and monitoring individuals. By contrast, the proposals in this report are focused on places, patterns, and service delivery—not on placing citizens under heightened scrutiny.

The strategies discussed in this report can be distinguished from two high-profile examples of controversial, individual-focused predictive-policing programs: Chicago's Strategic Subject List (SSL) and Los Angeles's Operation LASER. Chicago launched SSL

in 2012 to identify individuals deemed at elevated risk of involvement in a shooting. The program ultimately expanded to more than 400,000 names before being discontinued in 2019 amid concerns about bias. Similarly, Los Angeles’s Operation LASER assigned “chronic offender” scores to individuals based on criminal history and police contacts. Those identified by the system were subjected to police attention, including warning letters, home visits, and department-wide dissemination of their personal information. Following an inspector-general audit that identified potential shortcomings, the program was terminated in 2019.

Regardless of the merits of these programs or the concerns they raised, these cases are instructive precisely because they illustrate what this proposal is not. SSL and LASER involved an effort to predict the behavior of specific individuals and then direct police resources toward them. Whatever one thinks of those programs, they involved a fundamentally different use of data from the one contemplated here. This proposal does not seek to score, rank, monitor, or target *people*. Instead, it uses aggregate information to identify geographic patterns, operational bottlenecks, and areas where public services can be deployed more effectively. The distinction between managing places and monitoring people provides an important civil liberties safeguard.

With that distinction in mind, Singapore offers a more constructive model. As part of its Smart Nation initiative, Singapore has employed AI to deploy resources, automate public services, reduce traffic, and more. In some cases, Singapore does use facial recognition technology to improve crime prevention. However, in the case of traffic, for example, the city primarily uses place-based environmental analytics rather than individual profiling. By analyzing macro data, camera feeds, and spatial patterns, its model monitors the flow of vehicles and transit to update schedules and traffic lights accordingly.²⁴

Predicting where drag racing is likely to occur based on traffic-camera data and road geometry is not the same as flagging individuals for enhanced scrutiny. Identifying which blocks are most likely to generate sanitation overflow based on permit patterns and collection schedules does not involve personal data at all. Modeling where nightlife disorder tends to cluster based on liquor-license density, street width, and late-night transit access involves no individual profiling whatsoever. The unit of analysis is a neighborhood, a corridor, a building, or a permit type—not a person.

This distinction is the difference between a system that helps a sanitation supervisor route trucks more efficiently and one that generates a list of people to monitor. Despite the debate about the latter, the city should readily embrace the former without hesitation.

Work Requirement Fix

A significant legal question is whether the city is allowed to implement AI automation in the first place. Local Law 144 of 2021—the Automated Employment Decision Tool (AEDT) law—took effect on January 1, 2023, with enforcement beginning that July after the Department of Consumer and Worker Protection adopted its final implementing rules.²⁵ The law prohibits employers and employment agencies from using automated tools to substantially assist hiring or promotion decisions unless the tool has undergone an independent bias audit within the previous year, the results of that audit are published, and candidates are notified that such a tool is being used to evaluate them. At the state level, Governor Kathy Hochul introduced the LOADinG Act, which requires the

disclosure of automated decision-making tools and prohibits the use of AI that “impact[s] the rights, civil liberties, safety, or welfare of an individual unless such utilization is authorized in law.”²⁶

At first glance, the AEDT law and LOADinG Act might appear to constrain the proposal outlined in this report. They do not. Their restrictions are relevant only to a computational process that issues a simplified output to substantially assist or replace discretionary decision-making for decisions that affect *natural persons*. The systems proposed here, however, score buildings, corridors, intersections, and permit types. They do not score job candidates or discriminate against individuals, and they make no employment decisions whatsoever. The same distinction that resolves the privacy question—managing places rather than monitoring people—resolves this question as well.

How the Mayor Makes a Difference

Despite not running afoul of Local Law 144, these innovations may face significant pushback, given anti-AI public opinion and union pressure. Here, the role of the mayor is crucial. Throughout every data revolution in this city’s history, executive leadership was crucial. Mike Flowers explained that, as head of the city’s first data analytics office, “None of this ever could have happened without Bloomberg’s backing.... Without mayoral support it’s just impossible.”²⁷ Bloomberg’s leadership and support for his “Geek Squad” team of data analysts were necessary for their success.²⁸

The mayor’s office, by itself, can spearhead significant reform because almost none of the proposals in this brief requires any new legislation. MODA already exists and has been codified in the city charter since 2018. The Open Data Law is already on the books. OTI already controls the data infrastructure. The mayor can direct the automation of the remaining datasets, order the modernization of MMR, commission the internal model described above, and convene the relevant commissioners—all by executive action. Just a few months ago, Governor Hochul signed an executive order launching what she called the “Regulatory Reset,” an effort that will use AI to analyze 18 million words of codified law to modernize obsolete regulations. Hochul said that the initiative will make government more effective, with AI serving only to flag potential issues; all final decisions will remain in human hands.²⁹ Likewise, the Commission on Government Efficiency, created precisely to modernize government and improve service delivery, is the natural vehicle for this agenda.

Personnel, as always, is policy. Bloomberg’s analytics revolution is traceable to a handful of appointments, Flowers’s chief among them. Mamdani’s most consequential technology decision will likewise be his choice of chief technology officer (CTO)—who will ideally have a record of actually building teams and products, not just managing vendor contracts. A coalition of civic technologists has already urged exactly this, calling for a “CTO who can build” supported by a dedicated digital service team.³⁰

Finally, mayoral attention is itself a forcing function. CompStat reduced crime not only because the maps were impressive but because precinct commanders had to stand before their superiors every week and answer for the numbers. The AI-era equivalent is straightforward: a recurring meeting at which agency heads answer for the quality-of-life indicators that their dashboards now surface in real time. A model can flag the corridor

where nightlife disorder is building or the block where trash overflow is imminent, but only the mayor can make commissioners accountable for acting on that information. While the technology supplies the foresight, the executive enforces the consequences.

How Building Makes a Difference

The final question is: Who should build all this? Typically, the default answer would be outside vendors. But that would be a mistake, as demonstrated by the city's recent track record. An analysis of Checkbook NYC spending data by the Tech Mayor Project found that, in FY 2025 alone, New York spent at least \$769 million on outside digital services that could have been delivered in-house, with over \$2.2 billion flowing through vendors providing some mix of digital services and IT hardware.³¹ One Department of Buildings portal has consumed \$172.8 million since 2016, almost entirely for a single vendor.³² The CityTime payroll system became a several-hundred-million-dollar scandal that had to be settled in 2012.

The cautionary tale of the current era is MyCity. The Adams administration's flagship technology project was supposed to be a one-stop portal for all city services. Instead, contractors invoiced the city over \$100 million over two and a half years—spread across at least 43 private firms, many selected through the city's least transparent procurement methods—and produced a small fraction of what was promised.³³ Memorably, the project involved a small-business chatbot that dispensed illegal advice in its first weeks before having to be shut down. Worst of all, the underlying technology built by contractors is proprietary and not owned by the city, meaning that the city cannot fix, extend, or even fully inspect what it paid for. Every safeguard proposed earlier in this report, from immediate rollback to full data control, depends on ownership.

Other governments have shown the benefits of an alternative approach. The U.K.'s Government Digital Service, founded in 2011, consolidated hundreds of departmental websites into a single GOV.UK platform and saved the British government billions of pounds by building in-house.³⁴ The U.S. Digital Service, created in 2014 after in-house engineers rescued the failed HealthCare.gov launch, saved \$285 million in projected savings over the next five years,³⁵ with measurable improvements to everything from veterans' benefits to tax filing.³⁶ New Jersey's Office of Innovation, which uses predictive algorithms to improve state services, helped modernize New Jersey's outdated systems while delivering measurable results for a far lower cost than outsourcing.³⁷ In New York City, this would look similar to ACCESS NYC, a social-services eligibility website built originally by Bloomberg's team at DoITT in 2006.

The internal AI infrastructure described above—a self-hosted foundation model, fine-tuned on city data and governed through city-controlled access protocols—should be built, operated, and owned by an in-house NYC digital service team. This would be a blended group of engineers, data scientists, designers, and agency subject-matter experts working in a hub-and-spoke model across departments.

Part of the internal construction will require building Model Context Protocol (MCP) servers that act as a bridge between agentic LLMs and government data. The Free Law Project and BetaNYC have already created open-source MCP servers for the city's charter, legislative agenda, and history, as well as the comptroller's online budget tool, Checkbook NYC.³⁸ Rather than requiring an AI agent to interact separately with each dataset, MCP



provides a standardized way for agents to discover, access, and work with information from multiple data sources, making government data more accessible and useful for AI-assisted analysis.

The case for in-house ownership is not only about control but also performance. An internal AI reduces the number of intermediaries between operational need and technical response. For example, in 2012 the MTA contracted with outside vendors to create an NYC transit phone application. After spending thousands of dollars, to little effect, MTA pivoted to in-house construction. As Chief Customer Officer Shanifah Rieara explained, using an external vendor meant greater friction to innovation, updates, and service: “Having that capacity in-house means we can build it, maintain it and update it. As we get feedback from customers, it gets integrated by our team. We’re not shelling out hundreds of thousands of dollars,” said Rieara.³⁹ Building internally means fewer handoffs, fewer translation layers between agencies and vendors, and faster iteration cycles when systems need to be adjusted in real time.

Ownership also materially improves security and auditability. A city-hosted model can be constrained within municipal access controls, logged end-to-end, and kept within a defined perimeter rather than distributed across third-party environments. That containment reduces exposure points and simplifies oversight, while making it easier to enforce consistent data-governance standards across agencies.

Finally, customization becomes far more granular. AI systems are flexible tools that are improved constantly with experience, more data and updated objectives. When owned internally, they can be shaped by the direct “invisible hand” of inspectors, dispatchers, and field supervisors who use them daily, rather than being constrained by the fixed feature sets and update cycles of outside vendors. This is crucial in the world of AI, where updated models can radically improve capabilities in a short time.

For these reasons, predictive and generative systems should not be deliverables handed over at the end of a procurement cycle. They are a living infrastructure that requires continuous refinement, tight feedback loops, and institutional knowledge embedded directly inside the team that builds them.

Endnotes

- ¹ New York City Office of the Mayor, “Mayor Mamdani Announces Commission on Government Efficiency to Make Government Work Better for New Yorkers,” press release, May 28, 2026.
- ² Rafael A. Mangual, “New York’s Uneven Crime Drop,” *City Journal Substack* (blog), Aug. 26, 2026.
- ³ New York City Police Department, *Citywide Quality of Life Statistics* 2, no. 35, report covering Aug. 24–30, 2026; Rich Calder and David Spector, “Public Urination, Drug Use, Other NYC Quality-of-Life Complaints Are Surging and Will Only Get Worse Under Mamdani: Critics,” *New York Post*, Dec. 13, 2025.
- ⁴ Ana Champeny et al., “Straight from New Yorkers 2025: Residents Still Largely Dissatisfied with Quality of Life, Services, and Safety,” Citizens Budget Commission, June 4, 2025.
- ⁵ Office of the Governor of New York, “Murders in NYC, 1993–2025,” June 2025.
- ⁶ Hyunseok Jang, Larry T. Hoover, and Hee-Jong Joo, “An Evaluation of Compstat’s Effect on Crime: The Fort Worth Experience,” *Police Quarterly* 13, no. 4 (2010): 387–412.
- ⁷ Franklin E. Zimring, *The City That Became Safe: New York’s Lessons for Urban Crime and Its Control* (New York: Oxford University Press, 2011).
- ⁸ Adam Groff, “CompStat,” EBSCO Research Starters, 2013.
- ⁹ Randy Rieland, “How Data and a Good Algorithm Can Help Predict Where Fires Will Start,” *Smithsonian Magazine*, Mar. 2, 2015.
- ¹⁰ Nicole Gelinas, “The Fourth Urban Revolution,” *City Journal*, Spring 2016.
- ¹¹ City of New York, “NYC Open Data.”
- ¹² See Stop the Chop NY/NJ, <https://stopthechopnynj.org>.
- ¹³ N.Y.C. Local Law No. 11 (2012).
- ¹⁴ New York City Office of the Mayor, “Open Data for All,” July 2015.
- ¹⁵ City of New York, “NYC Open Data,” Dashboard.
- ¹⁶ Rachael Fauss, testimony on behalf of Reinvent Albany before the New York City Council Committee on Technology, *Oversight—Open Data Compliance*, Feb. 27, 2024.
- ¹⁷ NYC Open Data, “Feasibly Automating Your Data,” April 2018.



- 18 David Tussey and Jun Yan, “Principles for Open Data Curation: A Case Study with the New York City 311 Service Request Data,” *arXiv* preprint (revised March 2025): 2502.08649.
- 19 Julia Nagle, “5 Ways to Improve New York City’s Mayor’s Management Report,” Citizens Budget Commission, Oct. 8, 2024.
- 20 New York City Office of the Mayor, “Preliminary Mayor’s Management Report,” March 2026, 23.
- 21 See <https://www.nyc.gov/site/operations/research/mayor-office-of-data-analytics.page>
- 22 City of New York, *Inter-Agency Data Exchange Agreement*, 2010.
- 23 Emma Neuwirth, “Launching an Artificial Intelligence Powered Virtual Assistant for Commerce Center Reporting,” J.P.Morgan Payments, Oct. 29, 2025.
- 24 Alea IT Solutions, “Singapore Introduces AI-Powered Traffic Management System: A Step Towards Smart City Success,” Oct. 7, 2024.
- 25 Simone R. D. Francis et al., “New York City Adopts Final Rules on Automated Decision-Making Tools, AI in Hiring,” *National Law Journal*, Apr. 8, 2023.
- 26 Legislative Oversight of Automated Decision-Making in Government Act (LOADinG Act), New York Senate Bill S7543A (2023–24).
- 27 Mike Flowers, interview by Eddie Copeland, “NYC’s Mike Flowers Speaks to Policy Exchange About Data Driven Cities,” LinkedIn, Oct. 27, 2015.
- 28 Alan Feuer, “The Mayor’s Geek Squad,” *New York Times*, Mar. 23, 2013.
- 29 Grace Ashford, “Fax and Telegram? New York’s Archaic Requirements Get an A.I. Overhaul,” *New York Times*, July 8, 2026.
- 30 See <https://web.archive.org/web/20260412130140/https://techmayor.nyc/p/dear-zohran>
- 31 See <https://web.archive.org/web/20260112201126/https://techmayor.nyc/p/nyc-pays-vendors-billions-for-government-tech>
- 32 New York City Council, *Report on the Fiscal 2027 Preliminary Plan, the Fiscal 2027 Preliminary Capital Commitment Plan, and the Fiscal 2026 Preliminary Mayor’s Management Report for the Committee on Housing and Buildings* (March 2026).
- 33 Zachary Groz, “How Eric Adams’s MyCity Portal Became a \$100 Million Question Mark,” *New York Focus*, Mar. 19, 2025.
- 34 “Government Digital Service,” Wikipedia, last modified Sept. 6, 2026.
- 35 U.S. Digital Service, “USDS by the Numbers,” *2024 Impact Report*, 2024.
- 36 U.S. Digital Service, *2024 Impact Report*, 2024.



- ³⁷ Colin Wood, “Under a New Governor, New Jersey Will Keep Its Innovation Office and Its Innovation Chief,” *State Scoop*, Jan. 5, 2026.
- ³⁸ Noel Hidalgo, “New Digital Democracy Tools for AI—MCPs for the Charter/Laws, Council, the City Record, and Checkbook,” BetaNYC, May 21, 2026.
- ³⁹ Shanifah Rieara, Kevin Call, and Jen Chen, interviewed by Nicole Gelinas, “How the MTA Built Its New App In-House,” *Vital City*, Apr. 22, 2026.