



BEN BARTLETT
CITY COUNCILMEMBER, DISTRICT 3

IGOR TREGUB
COUNCILMEMBER, DISTRICT 4

08

CONSENT CALENDAR

March 10, 2026

To: Honorable Mayor and Members of the City Council
From: Councilmember Ben Bartlett (Author) and Councilmember Igor Tregub (Co-Author)
Subject: The Berkeley Rule: Artificial Intelligence Municipal Framework

RECOMMENDATION

1. Adopt a resolution which affirms Berkeley’s commitment to innovation in service of liberty, dignity, and the public good.
2. The City hereby endorses “The Berkeley Rule” and encourages City staff, commissions, and community partners to consider its Ten Principles in the planning, evaluation, and oversight of artificial intelligence systems. The Berkeley Rule: Put Residents First; Modernize City Services; Empower the Community; Ensure Transparency and Accountability; Standardize Operations; Certify Ethical Use; Protect and Prepare Our Workforce; Defend Civil Liberties; Social Advancement and Accessibility; and Catalyze Civic Wealth. The principles are meant to encourage the Artificial Intelligence (AI) industry to develop products and services aligned with these goals.

POLICY COMMITTEE RECOMMENDATION

On January 28, 2026, the Health, Life Enrichment, Equity & Community Committee adopted the following action: M/S/C (Taplin/Tregub) to send the item to Council with a positive recommendation, and that items 3 and 4 in the recommendation be amended to be less prescriptive.. Vote: All Ayes.

BACKGROUND

The Berkeley Rule serves as a mental model for City staff, guiding effective, ethical AI deployment in municipal operations. This framework also signals to the private sector the AI tools and innovations Berkeley seeks to procure and pilot for municipal challenges.

Artificial Intelligence (AI) refers to computer systems and robotics technologies, including autonomous or semi-autonomous machines, that perform tasks typically requiring human intelligence, movement, or judgment. This includes, but is not limited to: recognizing speech, interpreting images, decision-support

systems, independent analysis of information, generating text and images, and performing physical tasks through autonomous movement, manipulation, or interaction with the environment.

More specifically, AI systems use large datasets and advanced algorithms to identify patterns, make predictions, or generate content based on input data. Generative AI, a rapidly growing family of AI models, powers tools such as ChatGPT, Claude, Llama, and Gemini, which generates human-like text, images, or code. AI is having widespread impact across sectors. In government, it is being used to streamline service delivery, detect fraud, assist in emergency response, and analyze infrastructure maintenance needs. According to a 2023 report by the McKinsey Global Institute, generative AI could add up to \$4.4 trillion annually to the global economy, while also reshaping labor markets and intensifying debates around fairness, transparency, and data privacy.¹ In the public sector, cities like New York, San Francisco, and Washington, D.C. are beginning to implement AI use policies that emphasize accountability, human oversight, and equity.²

The adoption of AI technologies is rapidly expanding, with an increasing number of individuals incorporating AI into their daily lives. According to a 2024 report by Statista, nearly 40% of Americans reported using AI-powered tools, such as virtual assistants, search engines, or recommendation systems, daily.³ Furthermore, a 2024 survey by McKinsey found that 65% of businesses are already utilizing AI in some capacity, with significant integration into customer service, data analysis, and process automation.⁴ These figures indicate that AI is not only being widely adopted by consumers but also becoming integral to various sectors, highlighting its broad utility and growing presence. Additionally, the AI services market is projected to reach \$243 billion by 2025, highlighting the increasing reliance on AI across industries.⁵ A significant 25% of enterprises are expected to deploy AI agents this year, demonstrating the growing adoption of AI-driven solutions to improve efficiency and decision-making.⁶ McKinsey's 2023 report reveals that nearly half (49%) of tech leaders now say AI is fully integrated into their business

¹ McKinsey & Company. *The Economic Potential of Generative AI: The Next Productivity Frontier*. June 2023.

<https://www.mckinsey.com/business-functions/mckinsey-digital/our-insights/the-economic-potential-of-generative-ai-the-next-productivity-frontier>.

² City of New York. *AI Action Plan*. 2023 <https://www.nyc.gov/assets/oti/downloads/pdf/reports/artificial-intelligence-action-plan.pdf>

³ Statista. (2024). Percentage of U.S. population using AI tools regularly. <https://www.statista.com/forecasts/1480449/ai-tools-popularity-share-usa-adults>

⁴ McKinsey & Company. (2024). State of AI in Business <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-2024>

⁵ Marr, B. (2025, March 10). 15 Mind-Blowing AI Statistics Everyone Must Know About Now. Forbes.

<https://www.forbes.com/sites/bernardmarr/2025/03/10/15-mind-blowing-ai-statistics-everyone-must-know-about-now/>

⁶ Marr, B. (2025, March 10). 15 Mind-Blowing AI Statistics Everyone Must Know About Now. Forbes.

<https://www.forbes.com/sites/bernardmarr/2025/03/10/15-mind-blowing-ai-statistics-everyone-must-know-about-now/>

strategy, a clear indication of its essential role in modern organizational operations.⁷ This widespread integration reflects the remarkable increases in AI usage, with enterprises harnessing its potential to streamline operations, enhance customer experiences, and drive growth.

The global AI market is projected to grow from \$208 billion in 2023 to \$1.85 trillion by 2030,⁸ and over 65% of organizations worldwide are expected to adopt AI by 2024, a significant increase from just 20% in 2017.⁹ Cities are already utilizing AI technologies in areas such as law enforcement, traffic management, and tenant screening, with over 40 major U.S. cities employing tools like predictive policing, automated license plate readers, and AI-powered chatbots, often without adequate oversight.¹⁰ While the White House Blueprint for an AI Bill of Rights and NIST's AI Risk Management Framework offer valuable guidance on AI governance, neither provides enforcement mechanisms.¹¹ Additionally, over 45 states are now drafting or debating AI-specific legislation, signaling the growing need for formal regulation.¹²

As AI technology rapidly evolves, public and private institutions continue to develop coherent policies and frameworks. At the same time, public sentiment reveals growing concerns that aren't always reflected in the integration and use of AI. For example, in the YouGov¹³ poll on AI regulation, many U.S. citizens still believe that there should be more regulation of AI. Accordingly, policy should meet the public's demand for proactive, growth-oriented regulation, while avoiding federal pre-emption.

⁷ McKinsey & Company. (2023). The State of AI. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai>

⁸ Artificial Intelligence - Worldwide <https://www.statista.com/outlook/tmo/artificial-intelligence/worldwide>

⁹ McKinsey & Company. (2024). The State of AI. <https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai-2024>

¹⁰ Brookings Institution (2023). The Geography of AI: Which Cities Will Drive the Artificial Intelligence Revolution? <https://www.scirp.org/reference/referencespapers?referenceid=3933815>

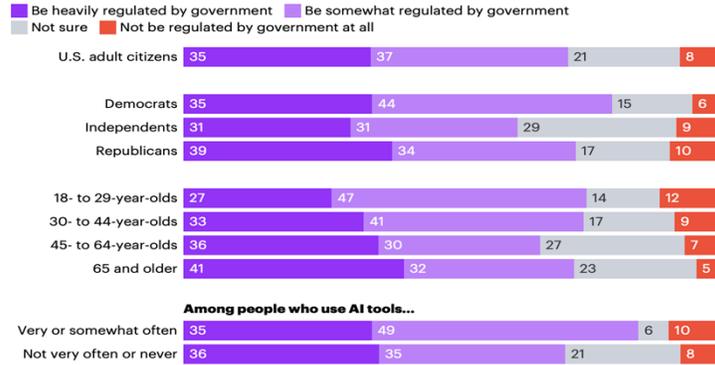
¹¹ The White House. (2022). Blueprint for an AI Bill of Rights <https://bidenwhitehouse.archives.gov/ostp/ai-bill-of-rights/>

¹² National Conference of State Legislatures. (2023). State Artificial Intelligence Legislation. <https://www.ncsl.org/technology-and-communication/artificial-intelligence-2024-legislation>

¹³ Orth, Taylor. "Americans Are Divided on AI's Societal Impact, but Most Support Government Regulation." *YouGov*, May 25, 2023 <https://today.yougov.com/politics/articles/45747-americans-are-divided-artificial-intelligence-poll>

Most Americans support government regulation of AI

Do you think that artificial intelligence (AI) should...? (%)



YouGov

The Economist/YouGov | May 20 - 23, 2023 14

AI Alignment

AI alignment is the process of incorporating human values and public goals into artificial intelligence systems to ensure they operate safely, ethically, and as intended.¹⁵ Alignment helps mitigate unintended consequences, ensuring that AI systems operate as intended and are consistent with human values and goals. For example, if one were to ask a generative AI chatbot how to build a weapon, it may either respond with instructions or refuse to provide potentially dangerous information. Unlike older logic-based AI and software approaches whose responses are manually coded by human programmers, a modern machine learning based AI model's response is determined by how the creators arranged it. In addition, modern AI systems learn from vast amounts of data based on how people behave online. That data reflects real human beliefs and habits, including many conscious and unconscious biases. As a result, today's AI systems can only be guided in general terms rather than precisely aligned with human values. While using human-like language can help us understand how AI systems work, it may also lead to distorted notions about AI's capabilities.¹⁶

Any municipal AI framework should incorporate principles of alignment and transparency. A major challenge is that, since nearly every principle conflicts with others, nontrivial processes are needed to resolve contradictions that confront AI Alignment.¹⁷ Deployments of AI systems should specify such processes.

¹⁴ Orth, Taylor. "Americans Are Divided on AI's Societal Impact, but Most Support Government Regulation." *YouGov*, May 25, 2023. <https://today.yougov.com/politics/articles/45747-americans-are-divided-artificial-intelligence-poll>

¹⁵ Jonker, Alexandra, and Alice Gomstyn. "What Is AI Alignment?" IBM, April 17, 2025. <https://www.ibm.com/think/topics/ai-alignment>.

¹⁶ De Kai. *Raising AI: An Essential Guide to Parenting Our Future*. 2025. MIT Press. <https://dek.ai/raising-ai>.

¹⁷ De Kai. "Should A.I. Accelerate? Decelerate? The Answer Is Both." *New York Times*, December 10 2023. <https://www.nytimes.com/2023/12/10/opinion/openai-silicon-valley-superalignment.html>

Aligning AI with Berkeley's Values & Strategic Goals

When aligned with Berkeley values, AI can help advance the city's strategic goals, including social advancement, improved public services, environmental protection, and civic trust. As Berkeley adopts new technologies, AI could be used to support, not replace, human judgment, expand access to essential services, and accelerate progress on the priorities most important to residents. The following sections describe how AI can support the values and goals set forth in the City of Berkeley Strategic Plan:

1) **Improve residents' lives by delivering accessible and innovative services.**

AI presents the opportunity to liberate residents from bureaucratic friction and unnecessary expense. When thoughtfully implemented, AI can streamline processes, support both residents and staff, and promote fairer access to public resources. The result is a more efficient, responsive, and inclusive city that better serves the community.

2) **Safeguarding civil liberties, equity, and democratic participation.**

Berkeley's commitment to civil liberties and democratic governance requires that any use of AI protects privacy, free expression, and due process. As AI is integrated into City services, it must align with Berkeley's civil liberties framework and strengthen the protection of fundamental rights.

AI systems should be transparent, include public input, and remain subject to human oversight. When guided by equity and democratic accountability, AI can support fair employment practices, workforce development, and inclusive economic growth while safeguarding labor rights.

With strong privacy protections and meaningful community engagement, AI can reduce bureaucratic barriers, expand access to public resources, and advance Berkeley's commitments to justice, sustainability, and democracy.

3) **Create affordable housing and support services for our most vulnerable community members.¹⁸**

¹⁸City of Berkeley, *Strategic Plan* (Berkeley, CA: City of Berkeley), accessed July 10, 2025, <https://berkeleyca.gov/your-government/our-work/strategic-plan>.

AI-aligned decision tools can help Berkeley distribute housing assistance more equitably and improve planning for affordable housing.^{19,20} These systems could identify high-need areas, streamline application processes, and inform land-use decisions, making housing support more accessible and effective. AI can also support the City's efforts to produce, preserve, and protect housing by improving financing strategies, prioritizing maintenance, and targeting subsidies to prevent displacement and preserve existing units.

For example, this approach can directly support Berkeley's Middle Housing ordinance by enabling more efficient planning for duplexes, triplexes, and fourplexes across the city.²¹ By expanding housing options and improving affordability, aligned AI tools can help advance Berkeley's commitment to housing security for its most vulnerable residents.²²

4) Be a global leader in addressing climate change, advancing environmental justice, and protecting the environment.

AI alignment can help advance the city's climate sustainability plans and goals such as the Berkeley's Climate Action Plan and Realize Vision 2050²³ by improving infrastructure condition assessments, emissions tracking, energy optimization, and environmental justice analysis. These tools can support smarter infrastructure planning, including water reuse systems, grid modernization, and targeted decarbonization efforts.²⁴

AI can also strengthen climate resilience through monitoring of heat, air quality, and wildfire risk,²⁵ while supporting urban greening, irrigation management, solar and heat pump permitting, and enforcement in underserved neighborhoods.²⁶ By directing investments where they have the

¹⁹ Orozco, D., & Das, S. (2023). *How AI Can Help Cities Advance Housing Equity*. Brookings Institution.

<https://www.brookings.edu/articles/building-ai-cities-how-to-spread-the-benefits-of-an-emerging-technology-across-more-of-america/>

²⁰ University of California, Berkeley Urban Displacement Project. (2022). *Using Data and Technology to Advance Housing Justice*.

<https://www.urbandisplacement.org/>

²¹ City of Berkeley. (2024). *Middle Housing Ordinance Overview*. <https://berkeleyca.gov/construction-development/land-use-development/general-plan-and-area-plans/middle-housing-zoning#:~:text=Project%20overview,in%20a%20range%20of%20sizes.>

²² California Department of Housing and Community Development. (2023). *Affirmatively Furthering Fair Housing in Local Planning* <https://www.hcd.ca.gov/planning-and-community-development/affirmatively-furthering-fair-housing>

²³ [https://berkeleyca.gov/sites/default/files/documents/2025-12-](https://berkeleyca.gov/sites/default/files/documents/2025-12-02%20Special%20Item%2001%20Sharing%20the%20Final%20Report%20of%20the%20Mayor.pdf)

[02%20Special%20Item%2001%20Sharing%20the%20Final%20Report%20of%20the%20Mayor.pdf](https://berkeleyca.gov/sites/default/files/documents/2025-12-02%20Special%20Item%2001%20Sharing%20the%20Final%20Report%20of%20the%20Mayor.pdf)

²⁴ World Economic Forum. (2022). *Harnessing Artificial Intelligence for the Earth*.

https://www3.weforum.org/docs/Harnessing_Artificial_Intelligence_for_the_Earth_report_2018.pdf

²⁵ National Institutes of Health (2023). *Smart City Tools for Air Quality and Heat Monitoring*

<https://pmc.ncbi.nlm.nih.gov/articles/PMC10280551/>

²⁶ Climate TRACE. (2023). *AI-Based Emissions Tracking Platform*.

<https://climatetrace.org/explore/#admin=&gas=co2e&year=2024&timeframe=100§or=&asset=>

greatest impact, AI can help Berkeley reduce emissions, optimize energy use, and promote environmental justice for a healthier, more sustainable future.

5) **Champion Social Advancement.**²⁷

AI aligned with Berkeley’s social advancement goals can help strengthen labor protections, uphold civil rights, and prioritize community needs across health, housing, and employment programs. By monitoring for disparate impacts and incorporating ongoing feedback from historically marginalized communities,²⁸, these systems can support equity, social mobility, and resident well-being.

AI can also expand public participation by making local democracy more accessible and understandable. Tools such as plain-language summaries, multilingual public comment translation, and participatory budgeting simulations can help residents engage more meaningfully in City decision-making^{29,30}. Together, these uses of AI can make local government more transparent, inclusive, and responsive to the community.

6) **Provide an efficient and financially-healthy City government.**³¹

AI can help Berkeley shift away from regressive revenue models that rely on fines and fees and toward support, prevention, and equity. AI audits can identify the City’s reliance on penalties that disproportionately impact low-income residents,³², while predictive tools can prevent fines through real-time reminders, hardship-based forgiveness, sliding-scale assessments, and service-based alternatives. This approach moves enforcement from punishment to restoration.

AI can also help the City maximize public assets and improve operational efficiency. By reusing proven models across departments, adopting open-source frameworks, and automating administrative oversight, Berkeley can reduce costs and increase transparency. AI tools can modernize lease and asset management by standardizing data, tracking city-owned properties in

²⁷ City of Berkeley, *Strategic Plan* (Berkeley, CA: City of Berkeley), accessed July 10, 2025, <https://berkeleyca.gov/your-government/our-work/strategic-plan>.

²⁸ City of Berkeley. (2021). *Health Equity and Racial Justice Framework*. <https://berkeleyca.gov>

²⁹ Introducing OpenAI for Government <https://openai.com/global-affairs/introducing-openai-for-government/>

³⁰ Mozilla Foundation. (2023). *Building Inclusive AI for Public Engagement*. <https://www.mozillafoundation.org/en/research/library/public-ai/>

³¹ City of Berkeley, *Strategic Plan* (Berkeley, CA: City of Berkeley), accessed July 10, 2025, <https://berkeleyca.gov/your-government/our-work/strategic-plan>.

³² Upturn. (2021). *The Use of Fines and Fees in U.S. Cities*. <https://www.upturn.org>

real time, and flagging inconsistencies or renewal risks, addressing long-standing gaps in oversight.

In addition, AI can support responsible revenue generation by identifying underutilized assets, improving compliance, and optimizing systems such as parking, business licensing, and short-term rental enforcement. By modernizing asset and revenue management, Berkeley can strengthen accountability, improve efficiency, and support long-term fiscal health while advancing its equity goals.

7) **Provide state-of-the-art, well-maintained infrastructure, amenities, and facilities.**³³

AI can help Berkeley monitor, maintain, and improve public infrastructure through predictive maintenance that identifies wear and tear in roads, water systems, lighting, and civic buildings. This approach allows the City to prioritize repairs, extend asset lifespans, and reduce costs and service disruptions.

AI can also improve transportation and mobility by reducing congestion, lowering emissions, and enhancing safety. Smart traffic systems can support Vision Zero goals, while accessible, real-time transit tools and AI-assisted public input can ensure community needs shape infrastructure decisions.

Within the education and youth service sectors, AI can help close opportunity gaps through personalized academic support, multilingual college and career guidance, and better targeting of resources. Together, these uses can help ensure Berkeley's public services are inclusive, efficient, and future-ready.

8) **Foster a dynamic, sustainable, and locally-based economy**³⁴

AI could strengthen Berkeley's economy by connecting residents to emerging job markets and supporting small businesses with accessible, data-driven tools. By analyzing labor trends, AI could guide job training programs, inform local hiring, and support the City's Economic Dashboards and procurement strategies. For job seekers, AI-powered tools like resume builders, interview simulators, and job-matching platforms (e.g., Jobscan, LinkedIn AI Coach) help level

³³ City of Berkeley, *Strategic Plan* (Berkeley, CA: City of Berkeley), accessed July 10, 2025, <https://berkeleyca.gov/your-government/our-work/strategic-plan>.

³⁴ City of Berkeley, *Strategic Plan* (Berkeley, CA: City of Berkeley), accessed July 10, 2025, <https://berkeleyca.gov/your-government/our-work/strategic-plan>.

the playing field, especially for historically under-resourced communities.³⁵ Entrepreneurs benefit from AI assistants that explain permitting, grants, and legal basics in clear terms, lowering barriers to small business ownership.³⁶ AI assistants can help draft small business setup documents, such as incorporation forms, operating agreements, and basic contracts, lowering barriers to entry and reducing start-up costs. In the workplace, large language models (LLMs) integrated with tools like Microsoft 365 and Slack automate tasks like summarizing meetings, drafting emails, or translating documents boosting productivity and job satisfaction³⁷.³⁸ In fields like law and healthcare, they support professionals with document review and compliance checks.³⁹ For contractors, AI-powered rendering tools can generate visualizations and design options, making project planning faster, clearer, and more affordable.

The City can make these AI tools available to the broader community through libraries, workforce training centers, and small business resource hubs, ensuring that all residents benefit from the efficiencies and opportunities AI provides.

9) **Create a resilient, safe, connected, and prepared City.**

Aligned AI systems could be essential for enhancing public safety, emergency response, and urban resilience in Berkeley. These tools could support wildfire risk mapping, aiding zoning decisions and enforcement of defensible space in hillside neighborhoods, where threats are most acute. AI also enhances community preparedness, facilitates hazard response coordination, and supports data-driven resilience planning, as outlined in the City's Resilience Strategy. AI could improve public safety outcomes without expanding surveillance. The use of real-time language translation and medical history alerts during emergency calls enable quicker and more equitable responses for non-English speakers and individuals with disabilities. AI-assisted acoustic systems could distinguish between fireworks and gunfire more accurately, thereby reducing false alarms and fostering community trust. Predictive models analyzing heat, crime, and public health data

³⁵ Koustas, D., et al. (2023). Barriers to Employment and AI Tools. NBER. <https://conference.nber.org/confer/2025/DTs25/farronato.pdf>

³⁶ Ghosh, R. (2022). AI and Small Business Resilience. Brookings. <https://www.brookings.edu/articles/hybrid-jobs-how-ai-is-rewriting-work-in-finance/>

³⁷ Binns, R., et al. (2023). AI in Knowledge Work. Oxford Internet Institute. <https://www.oii.ox.ac.uk/news-events/2023-the-year-ai-redefined-work-skills-and-the-future-of-employment/>

³⁸ California State Bar. (2024). Ethical AI Use in Legal Services. <https://www.calbar.ca.gov/Portals/0/documents/ethics/Generative-AI-Practical-Guidance.pdf>

³⁹ Simbo.ai, "The Role of Artificial Intelligence in Streamlining Document Review Processes for Legal Practitioners in 2024," *Simbo.ai Blog*, accessed July 14, 2025, <https://www.simbo.ai/blog/the-role-of-artificial-intelligence-in-streamlining-document-review-processes-for-legal-practitioners-in-2024-251528/>.

could direct emergency outreach or mental health crisis teams to areas of highest need, avoiding reliance on carceral tools. In health and social services, AI could extend the City’s reach to vulnerable residents. Tools could guide users through enrollment in programs like CalFresh, Medi-Cal, and city wellness initiatives using plain language and mobile-first design. Mental health chatbots could provide anonymous, culturally competent emotional support, encouraging early access to care. For frontline staff, AI could flag service gaps, track referrals, and monitor outcomes, creating a more coordinated and compassionate care system.

10) Be a customer-focused organization that provides excellent, timely, easily-accessible service and information to the community⁴⁰

Properly aligned AI can help Berkeley deliver faster, fairer, and more accessible public services. AI-powered virtual assistants can enhance 311 services, streamline workflows, and support public health communications by providing clear, real-time responses to resident inquiries around the clock, reducing wait times and service backlogs. AI dashboards can give City staff real-time insights to proactively address community needs, reinforcing Berkeley’s commitment to transparency and responsive governance. Tools such as permit chatbots, proactive alerts, and simplified dispatch systems can reduce administrative burdens for both residents and staff.

Aligned AI can also modernize ticketing and permitting by making traditionally punitive processes more transparent and equitable. AI tools can guide residents through contesting tickets by explaining violations, assisting with evidence submission, supporting multilingual appeals, and tracking cases in real time. In permitting, AI can help complete applications, flag errors, triage requests, and accelerate low-risk approvals. Together, these tools can improve accountability, reduce bias, and make everyday interactions with City Hall more efficient and just.

11) Attract and retain a talented and diverse City government workforce⁴¹

AI alignment can modernize Berkeley’s human resources systems by supporting inclusive hiring, reducing bias in recruitment, and strengthening employee training and advancement. AI tools can

⁴⁰City of Berkeley, *Strategic Plan* (Berkeley, CA: City of Berkeley), accessed July 10, 2025, <https://berkeleyca.gov/your-government/our-work/strategic-plan>.

⁴¹City of Berkeley, *Strategic Plan* (Berkeley, CA: City of Berkeley), accessed July 10, 2025, <https://berkeleyca.gov/your-government/our-work/strategic-plan>.

help forecast staffing needs, identify skill gaps, and recommend professional development pathways, supporting the City's goal of building a diverse and future-ready workforce.

Properly aligned AI should support workers rather than replace them, reinforcing Berkeley's commitment to fair employment practices and workforce development. Public-sector examples show how AI can help identify internal talent, reduce turnover, and monitor indicators of burnout or attrition to support employee well-being.

To ensure these benefits, AI systems should operate transparently and with ongoing oversight to enhance human decision-making. Emerging governance frameworks highlight the importance of fairness, explainability, and accountability when using AI in public-sector workforce management.

Potential Pitfalls of AI Systems

While holding significant promise, artificial intelligence also presents serious risks if not carefully governed. Without appropriate safeguards, AI systems could produce inaccurate or misleading outputs, often referred to as "hallucinations," that may result in flawed decisions in high-stakes areas such as housing, public benefits, policing, and financial catastrophe.⁴² AI tools could also reinforce systemic biases if trained on incomplete or discriminatory data, leading to unequal treatment or outcomes, especially for communities of color, low-income residents, and people with disabilities.⁴³ Moreover, without transparent processes, the use of AI in surveillance or decision-making could violate privacy rights, reduce public trust, and erode due process protections, particularly when algorithms are used to allocate resources, flag individuals, or influence enforcement actions.⁴⁴ These dangers are amplified when AI systems operate without public scrutiny or accountability mechanisms. To successfully realize the benefits of AI while minimizing harm, the City of Berkeley should adopt a formal AI framework.

Toward Artificial Intelligence Municipal Use Guidelines

Currently, the city lacks a consistent, citywide approach to AI deployment. Often, AI tools are independently applied, without a centralized inventory or standardized procurement and oversight protocols. This fragmented approach could result in operational vulnerabilities.

⁴²DigitalOcean. (2023). What Is AI Hallucination? <https://www.digitalocean.com/resources/articles/ai-hallucination>

⁴³U.S. Government Accountability Office. (2021). Artificial Intelligence: An Accountability Framework for Federal Agencies and Other Entities. <https://www.gao.gov/products/gao-21-519sp>

⁴⁴AI Now Institute. (2018). Algorithmic Accountability Policy Toolkit <https://ainowinstitute.org/publications/algorithmic-accountability-policy-toolkit>

To reap the benefits of AI while mitigating its risks, the City of Berkeley should develop an Artificial Intelligence Systems and Alignment framework that includes: Put Residents First; Modernize City Services; Empower Community; Ensure Transparency and Accountability; Standardize Operations; Certify Ethical Use; Protect and Prepare Our Workforce; Defend Civil Liberties; Social Advancement and Accessibility; and Catalyze Civic Wealth. Any framework should consider the elements herein referred to as the Berkeley Rule.

The Berkeley Rule:

I. Put Residents First

The City of Berkeley is committed to ensuring that any AI framework prioritizes the well-being of residents above all else. This policy aims to liberate residents from bureaucratic friction, eliminate unnecessary expense, and expand fair access to city services.

The City could seek to streamline internal operations, reduce operational costs, enhance the coordination and delivery of public services. Thoughtfully implemented, AI could improve workflows such as permitting, resident request routing, and document processing, resulting in faster and more consistent outcomes across departments. Cities like New York and San José are already seeing results: New York uses AI to prioritize housing inspections and assist with city service requests.⁴⁵ Likewise, San José uses predictive maintenance tools to identify infrastructure issues before they become costly emergencies.⁴⁶ More efficient permitting and service response times reduce delays and frustration for individuals and businesses. Improved infrastructure management powered by predictive tools means fewer service disruptions, better street conditions, and faster emergency responses. By enhancing coordination across departments, residents experience a more connected, responsive, and equitable government.

The City of Berkeley could use AI to rapidly review its municipal code to eliminate unnecessary reports and cut red tape. By automating routine administrative tasks, AI could allow Berkeley staff to focus on higher-impact work that encourages critical thinking, creativity, and direct public engagement. Moreover, AI-driven data analysis could also help departments identify gaps, target resources more efficiently, and support long-term planning in areas such as housing, workforce

⁴⁵ Mayor's Office of the Chief Technology Officer, City of New York. (2023). *AI Action Plan*. pg 2 <https://www.nyc.gov/assets/oti/downloads/pdf/reports/artificial-intelligence-action-plan.pdf>

⁴⁶ City of San José. (2023). *AI and Emerging Technology Strategy*. <https://www.sanjoseca.gov>

development, and public health by exploring cooperative care models that address current institutional insecurity. For example, this could include shared health coverage pools for freelancers and families, collective care planning for gig workers, wellness reward programs to encourage healthy habits, and neighborhood networks for exchanging caregiving services. The City should also consider recognizing and rewarding residents who contribute to improving city systems through civic dividends and benefits, thus ensuring that modernization directly supports the people who live and work in Berkeley.

The ultimate goal of AI adoption is to improve the lives of Berkeley residents. Any framework should be guided by the principle of improving service quality and efficiency, ultimately contributing to a higher quality of life for all community members and the eventual elimination of regressive fines, fees, and taxes.

II. **Modernize City Services**

The city should advocate for responsible use of AI in modernizing its municipal operations to improve efficiency, responsiveness, and service delivery. Potential use cases include 311 service triage, service kiosks, assistance with permits and licensing, emergency dispatch optimization, autonomous service delivery, maintenance, and transport, and emergency air deployments. AI systems could provide predictive solutions for vital infrastructure, such as roads, energy systems, and sidewalks, water systems, public health, and public buildings before routine problems become critical issues.

Berkeley could also utilize AI to develop and pilot new service models to make city services more accessible. For example, prepaid service savings programs might help residents secure discounted utilities, while optimized parking revenue systems and a community-owned broadband network could improve affordability and equity. Smart licensing processes and value-based pricing might guide development toward inclusivity. Additionally, predictive dashboards, better use of idle fleets, and AI-powered maintenance scheduling could streamline operations, reduce downtime, and make everyday services more reliable.

A key area for modernization is the city's permitting process. The City could explore AI-powered tools to pre-check construction and building permit applications for compliance with zoning and building codes. This is likely to provide immediate feedback to applicants, reduce the

potential for costly errors, and significantly decrease staff review time and backlogs. To implement this without direct cost to taxpayers, the City could adopt models like CivCheck Permitting AI in Seattle, where the permit applicants pay a small pre-screening fee directly to the vendor.

Finally, to ensure accountability and public trust, the City could mandate explicit documentation of human oversight protocols and backup measures for all AI systems used in time-sensitive or safety-critical situations. This includes determining when and how much human intervention should be encouraged, as well as ensuring that non-automated options are available where applicable.

III. **Empower the Community**

Community trust is the foundation of ethical AI governance. For Berkeley to harness the benefits of artificial intelligence while safeguarding civil rights, residents should have a direct voice in how these systems are adopted and used.

AI Advisory Board: If resources permit the City of Berkeley could consider establishing an advisory board composed of labor representatives, community leaders, civil rights and disability advocates, entrepreneurs, ethicists, technologists, and academic experts to ensure that AI is utilized in an ethical, equitable, and transparent manner.

Digital Ombudsman: To further strengthen accountability, the City could consider implementing a dedicated AI Ombudsman, an AI bot, to serve as a public-facing point of contact for questions, concerns, and complaints regarding municipal AI use. This solution would support algorithmic redress by helping residents understand how decisions are made by AI systems, and by facilitating rapid review, appeal, or correction when errors or harms occur.

AI Sandbox: Similarly, the City should consider creating an AI Sandbox program to pilot innovative tools in low-risk environments. These pilot programs would be reviewed by the AI Advisory Board, evaluated with public input, and include opt-in participation only. This approach ensures that new AI technologies are tested responsibly, with community involvement and safeguards in place prior to broader implementation.

IV. **Ensure Transparency and Accountability**

The City of Berkeley could consider maintaining a public AI use registry. This registry could present a clear and accessible listing of every AI tool being used by City departments. For each system, the registry would provide detailed information, including what the system does, what data it uses, who oversees it, and how residents can ask questions or challenge its outcomes.

This registry could be critical for systems that impact high-stakes processes such as permits, housing applications, benefits, and enforcement. All entries should be written in plain language to ensure accessibility and kept up to date as new tools are adopted. By implementing this approach, the City can enhance public accountability and make sure residents are informed partners in the use of the new technology.

Moreover, vendors should demonstrate maximum possible explainability in AI systems deployed in high-stakes areas. Explainable AI means systems are designed so that their outputs can be interpreted by experts and made understandable to the public. Explainable AI is subject to significant technical limits, including misleading, contradictory, unstable, mismatched, counterintuitive explanations, as well as the illusion of explainability which is logically unavoidable.⁴⁷ User misinterpretation and human factors can lead to explainability pitfalls. While acknowledging the limitations, this registry could help build trust in the AI systems by improving the transparency and comprehensibility of their decision-making processes.

Building on its AI registry, Berkeley could expand transparency by introducing consent portals where residents manage how their data is used. Algorithmic bias tracking might help identify and address systemic inequities, while municipal data exchanges and open licensing marketplaces could create clear, accountable systems for sharing public data. These measures would ensure that the value of information is managed with fairness, oversight, and full resident awareness, and economic inclusion in the deployment of automated decision-making systems across all City operations.

⁴⁷ De Kai. *Raising AI: An Essential Guide to Parenting Our Future*, chapter 16. 2025. MIT Press. <https://dek.ai/raising-ai>

V. **Standardize Operations**

Berkeley should consider strengthening its AI governance by centralizing inventory management, streamlining procurement, and applying uniform oversight procedures across all departments.

AI tools deployed in City operations should align with certain cybersecurity standards, including encryption in transit and at rest, regular audits, and protections against injection or tampering. To further strengthen privacy and civil liberties safeguards the city could employ:

- a. Privacy Impact Assessments (PIAs) could include a public comment period and could be reviewed by the AI Advisory Board before AI system deployment.
- b. Integrate the Resident Data Rights Charter, a clear process for residents to opt out of AI-processed data collection, modeled after the California Consumer Privacy Act (CCPA) principles.

To ensure strong compliance, the City could consider adopting the following operational standards:

1. Risk-Based Tiering Framework: The City should establish a risk-based tiering framework to classify AI systems based on their potential impact on residents. This framework, to be developed by the City Manager and reviewed by the AI Advisory Board, would include at a minimum:
 - a. Tier 1: Low-Risk Systems: AI tools that support internal administrative tasks and do not directly impact the public's rights or safety. Examples: meeting summarizers, internal project management software.
 - b. Tier 2: Medium-Risk Systems: AI tools that interact with the public in non-critical ways or assist City employees in making decisions with a low-to-moderate impact. Examples: public information chatbots, initial sorting tools for permit applications.
 - c. Tier 3: High-Risk Systems: AI systems that have a significant direct impact on residents' rights, safety, finances, or access to essential services, or that manage critical infrastructure. Examples: systems used in housing or benefits eligibility, law enforcement, or critical infrastructure management. The principles outlined in the Berkeley Rule would apply to all tiers, with specific procedural suggestions scaling with the level of risk.

2. National Institute of Standards and Technology (NIST) Guidelines: All AI systems should consider adhering to the NIST frameworks as the foundation for enforcement mechanisms.
 - a. External Applications: For constituent-facing use cases, the City should consider pursuing ISO27001 and/or ISO9001 certifications, ideally within 12 months of deployment. For applications involving personally identifiable information (PII), SOC2 compliance would be encouraged. Tools like Vanta or Workstreet can support these efforts.
 - i. ISO 27001 and ISO 9001 are internationally recognized standards that help organizations manage information security and quality
 - ii. ISO 27001 focuses on establishing an Information Security Management System (ISMS)
 - iii. ISO 9001 focuses on establishing a Quality Management System (QMS).
 - iv. SOC2, or System and Organization Controls 2, is a framework developed by the American Institute of Certified Public Accountants (AICPA) to assess and report on the controls of a service organization relevant to security, availability, processing integrity, confidentiality, and privacy.
 - b. Internal Applications: Policies should consider enforcing NIST standards with periodic internal audits. Vendors such as can provide cost-effective solutions to ensure compliance.
 - c. Regulatory Parallels: Identify existing policies and referrals that regulate the use of technologies, hard and software.
 - i. Example: Council's 2017 referral to develop a franchise application policy for robotic deliveries.⁴⁸
3. Procurement Standards: All external vendors, whether providing platform applications or custom development work for the City, would comply with these measures. A procurement framework should include a checklist to verify compliance or establish reasonable timelines for vendors to meet these guidelines. Standard contract templates

⁴⁸ Councilmember Ben Bartlett. (2017, December 19). *Personal Delivery Service Franchise Agreements*. City of Berkeley, District 3. <https://records.cityofberkeley.info/PublicAccess/api/Document/AY6Z5utuMHPCP6K72V2ETpqWhGYOZF7GÉhPtOJPDmkxowc51O buEKBjeoFzegHorlARtR4LfcRwbJpbwojXg3yY%3D/>

(MSAs) would include language for correction timelines in response to adverse events or audit findings.

4. Considerations for Return on Investment (ROI) Posture:
 - a. Direct ROI: Key performance indicators (KPIs) for AI deployments would include measurable benchmarks, such as time saved in administrative tasks, reduced turnaround times for policy implementation, and other quantifiable outcomes. Pilot programs should follow a structured framework to validate vendor claims.
 - b. Indirect ROI: AI initiatives should align with Berkeley’s core values, such as affordability, access, education, environmental sustainability, and social advancement. Vendors might justify how their solutions support these goals, ensuring no adverse impacts on these key pillars. Agentic LLM deployments would include guiding principles in their context windows to notify users and administrators of any misalignment with City values.
5. Possible Ongoing Governance and Evaluation: deployed systems could benefit from the following:
 - a. Annual Re-validation: The AI Advisory Board could conduct and publish an annual review of all Tier 3 (High-Risk) systems to re-validate their safety, performance, and equity impacts.
 - b. Incident Response: All issues reported to the Digital Ombudsman could be investigated and logged in the public AI Use Registry to further assure transparency. Significant incidents would be escalated to the AI Advisory Board for formal review and remediation.
6. System Decommissioning: A formal decommissioning plan could be encouraged for all Tier 3 (High-Risk) systems and specify the protocols for data processing, retention and disposal in compliance with City Auditor and public records requirements, and include a plan for the transition of public services to ensure continuity.
7. Government AI (GovAI) Coalition: the city should consider joining the Government AI (GovAI) Coalition facilitated by the City of San Jose⁴⁹ at no cost and partner with other cities to evaluate the adoption of its toolkits and practices.

⁴⁹ <https://www.sanjoseca.gov/your-government/departments-offices/information-technology/ai-reviews-algorithm-register/govai-coalition>

By implementing these suggested operational standards, compliance measures, and procurement accountability practices, Berkeley can ensure that AI systems are deployed responsibly, transparently, and in alignment with the City's strategic goals, while maximizing both direct and indirect benefits for its residents.

VI. **Certify Ethical Use**

The City of Berkeley could consider collaborating with not-for-profit organizations to establish independent AI ethics certification programs, ensuring safety, fairness, and accountability in its work with vendors who build and deploy artificial intelligence systems. These certifications should establish clear guidelines for safety, environmental responsibility, justice, and transparency. To achieve this, the City can partner with organizations with expertise in technology, policy, civil rights, and public interests to develop standards for its procurement procedures.

Certifications should include independent audits, inclusive design principles, and community impact protections to ensure that AI systems align with Berkeley's values. The City should prioritize workforce transition plans, explicit equitable targets, and methods to track performance and results over time in all AI-related contracts.

To further support ethical use, the City could establish a formal appeals mechanism for decisions made or influenced by AI systems. This mechanism would allow citizens to contest outcomes, request human review, and receive timely and accessible explanations.

By implementing an ethical use certification program and ensuring accountability measures, Berkeley should set a high standard for the responsible and transparent deployment of AI technologies.

VII. **Protect and Prepare Our Workforce**

As the City introduces more AI tools into its government operations, it should ensure that its workforce is supported and protected. Before rolling out a new system, departments could prepare a Workforce Impact Statement that examines how the technology may alter job duties, identifies any trainings, and outline opportunities for employees to transition into new roles.

These plans should be reviewed by a labor-management team that includes union representatives, with the goal of no layoffs resulting from AI adoption. Instead, AI should be utilized to automate repetitive tasks, freeing up time for public service and creating space for meaningful work.

The City should also invest in retraining and upskilling programs, ensuring that employees have access to continuing education, technical certifications, and cross-training opportunities to remain competitive and fulfilled in their roles. New professional development pathways should be created to help staff grow alongside advancing technology, reinforcing Berkeley's commitment to a strong, future-ready public workforce. During its implementation, to foster an inclusive approach to regulation development, staff engagement initiatives could be pursued through employee surveys and internal research to better understand their attitudes, expectations, and knowledge of AI threats and capabilities, while also empowering them to take the lead in developing case proposals.

As part of its ethical AI framework, the City could invest in public-sector innovation fellowships, rotational learning programs, and incentives for internal talent development. By protecting worker rights and preparing employees for the future of municipal service, Berkeley can lead in equitable workforce transformation.

VIII. **Defend Civil Liberties**

The City of Berkeley should safeguard civil liberties by not using facial recognition, biometric surveillance, or real-time tracking unless approved by the City Council, and then only after a robust public process is conducted and clear legal safeguards are in place. One potential use could be the use of a Privacy Impact Assessment (PIA). These AI systems assessments should not use sensitive or personal data, and any use thereafter should require a formal PIA to be completed before deployment. The privacy and civil liberties protections would identify what data is being used, how it is protected, and whether there are safer alternatives. The results of these assessments would be made public to ensure transparency and accountability.

The City could consider developing a set of Resident Data Rights to accentuate transparency and give people more control over their information in order to ensure it is as transparent as possible when AI is being used, how decisions are being made, and how to request human review or deletion of their data.

Furthermore, the City of Berkeley affirms that all AI systems with enforcement capabilities should preserve human judgment at the point of action. No automated system may take punitive or coercive measures, such as issuing citations, restricting access, or initiating legal consequences

without meaningful human oversight, and rapid appeal. Nor may any AI system employ martial force under any circumstance. These safeguards ensure that innovation serves to protect due process, community trust, and individual freedom.

IX. Social Advancement and Accessibility

AI developers could be encouraged to create a productized approach for Risk Assessment for high-impact AI technologies, particularly those used in housing, public safety, and transportation, and code enforcement to encourage AI systems to improve life for all residents. These evaluations should include pre-deployment social impact forecasts, conducted by third-party auditors using demographic data, scenario modeling, and participatory input to identify and mitigate potential harms. If these evaluations uncover biased outcomes, harmful system behavior, or deeper structural inequities, the City could take immediate action to correct both the technology and the underlying condition.

Berkeley can ensure that its AI systems promote fairness and inclusive outcomes by adopting standardized bias assessment frameworks. Tools developed by the Algorithmic Justice League and requirements such as New York City’s Local Law 144 can help measure, disclose, and mitigate disparate impacts across communities.

Berkeley could use AI to promote cultural growth and shared prosperity. With AI as an accessible tool for social advancement, residents might have opportunities to co-invest in local solar and renewable energy projects, crowdfund green infrastructure, and share in the benefits of climate impact returns. Housing equity programs and land stewardship trusts could expand access to stable homes. Meanwhile, cultural life would be free to flourish through community-funded art restoration, neighborhood business pools, festival revenue sharing, and pop-up market activations. By blending sustainability with cultural vitality, these initiatives could ensure that AI benefits every resident in our community.

X. Catalyze Civic Wealth

Cities everywhere struggle with budget deficits and shrinking revenue bases. Berkeley has the opportunity to move beyond this scarcity model. Beyond compliance and cost recovery, AI offers a transformational path to design new forms of civic wealth and municipal entrepreneurship. This approach opens pathways to shared value creation and community-owned innovation rather than relying on regressive fines, fees, or incremental efficiencies.

Artificial intelligence presents opportunities for the City of Berkeley, in partnership with nonprofit organizations, local innovators, and mission-aligned enterprises, to explore new ways of maximizing the value of public assets and advancing shared prosperity. Potential applications include identifying underutilized land, supporting dynamic leasing strategies, and forecasting value-based permitting opportunities responsive to changing economic conditions. Predictive tools may help surface untapped revenue potential, inform collaborative public-private and public-benefit partnerships, and support innovative models of municipal and community entrepreneurship.

Through such partnerships, AI-enabled systems could help activate vacant land, open rooftops for solar energy and urban farming, and optimize public facilities for broader community use. Digital billboards and archival collections may be responsibly licensed to generate cultural revenue, while vehicles and public equipment could be managed through shared-use models that return value to the community. Concepts such as community venture funds, real estate investment pools, and land value growth sharing illustrate ways residents might participate more directly in the city's prosperity. Emerging data-driven opportunities, including municipal data markets, neighborhood sentiment exchanges and digital twin licensing, further highlight Berkeley's potential role as a hub of civic and digital innovation.

More broadly, AI creates space to prototype new civic economies through cross-sector collaboration, where services are co-produced, benefits are equitably shared, and public data serves as a platform for innovation. Illustrative examples include decentralized licensing for local creators, micro-contracting opportunities for small businesses, and revenue-sharing models that support community infrastructure. In this context, public benefit-sharing approaches in civic technology collaborations reflect Berkeley's values, including shared intellectual property arrangements, royalty structures, open-source access, and reinvestment of proceeds into community-led initiatives and digital equity efforts.

In this vision, AI becomes a catalyst for inclusive prosperity and long-term fiscal resilience.

The City's commitment to its strategic plan would benefit from a high-level framework being codified into operational AI standards. The absence of a formal review process for algorithmic systems has allowed for multiple pathways of adoption, each with variable levels of risk, due process protection, and labor input. These Ten guidelines should serve as the foundation of Berkeley's AI Framework and be

embedded in the Berkeley Rules: Put Residents First; Modernize City Services; Empower Community; Ensure Transparency and Accountability; Standardize Operations; Certify Ethical Use; Protect and Prepare Our Workforce; Defend Civil Liberties; Social Advancement and Accessibility; and Catalyze Civic Wealth.

Comparative Civic Innovation Models

San Jose - Human-centered design principles guide the development of systems, and privacy safeguards protect sensitive information.⁵⁰ Security and safety are ensured through safeguards, and personnel empowerment is prioritized through education, training, and collaborative opportunities. These concepts should be linked to community benefits and human monitoring to ensure reliable AI deployment.

Boston - The 2023 “Interim Guidelines for Using Generative AI”⁵¹ from the City of Boston emphasize that public servants are still responsible for AI-generated material and offer a framework for responsible experimentation with programs like ChatGPT, Bard, and DALL·E. While cautioning against relying too heavily on unconfirmed AI results and encouraging the equitable, open, and safe use of these tools, the City encourages learning through workshops and provides contacts and resources for further research.⁵²

Seattle - The City of Seattle has announced its Generative Artificial Intelligence (AI) policy, which enables staff to utilize new technologies while adhering to established standards. The policy requires employees to obtain AI technology through approved procurement channels, review output to ensure consistency with City standards, attribute AI-generated content to the AI system, and ensure data is free of harmful bias, privacy concerns, and complies with the State of Washington Public Records Act and City policies. The guideline also requires a human to oversee the AI technology review.⁵³

Washington, D.C.- Washington, D.C.’s AI Values and Strategic Plan outlines⁵⁴ a citywide approach for the safe, equitable, and effective use of artificial intelligence in local government. The plan, based on five

⁵⁰ AI Reviews & Inventory | city of san José. Accessed June 30, 2025. <https://www.sanjoseca.gov/your-government/departments-offices/information-technology/digital-privacy/ai-reviews-algorithm-register>

⁵¹ City of Boston, *Interim Guidelines for Using Generative AI*, Version 1.1, prepared by Santiago Garces, Chief Information Officer, May 18, 2023. <https://www.boston.gov/sites/default/files/file/2023/05/Guidelines-for-Using-Generative-AI-2023.pdf>

⁵² City of Boston. *City of Boston Interim Guidelines for Using Generative AI*. Boston, MA: City of Boston, May 2023. <https://www.boston.gov/sites/default/files/file/2023/05/Guidelines-for-Using-Generative-AI-2023.pdf>

⁵³ “Responsible Artificial Intelligence (AI) Program.” Responsible Artificial Intelligence (AI) Program - Tech. Accessed June 30, 2025. <https://www.seattle.gov/tech/data-privacy/the-citys-responsible-use-of-artificial-intelligence>

⁵⁴ District of Columbia Office of the Chief Technology Officer, *DC’s AI Values and Strategic Plan*, 2023. <https://techplan.dc.gov/page/dcs-ai-values-and-strategic-plan>

guiding principles: transparency, accountability, justice, privacy, security, and inclusivity; emphasizes the wise application of AI to enhance public services while protecting the rights of residents.

Denver - The act mandates high-risk AI and system developers and deployers to protect the public from the risks of algorithmic discrimination. Developers must provide detailed disclosures, publicly summarize their systems, and notify the attorney general. Employers might implement risk management policies, conduct impact assessments, notify consumers, and provide mechanisms for data correction. Compliance with risk management framework provides an affirmative defense.⁵⁵

Chicago- Chicago's AI Principles⁵⁶, which place a strong emphasis on accountability, transparency, equity, dependability, privacy, and public involvement, provide a framework for the moral and efficient application of AI in local government. The city also places a high priority on public trust through interdisciplinary cooperation and participatory governance, emphasizing ongoing assessment and adapting AI systems in response to public input and real-world effects.

New York - The New York State Comprehensive guidelines for the responsible use of AI systems by state agencies, especially those that could have an immediate impact on the public, are established under NYS-P24-001: Acceptable Use of Artificial Intelligence Technologies. The regulation requires human oversight, which means that no entirely automated judgments that have a significant impact on the public are permitted. Instead, humans must continue to be held accountable and participate in the decision-making process. It encourages transparency by requiring public-facing systems to disclose their use of AI and places a strong emphasis on fairness by mandating authorities to monitor and correct bias. Agencies are required to maintain an AI inventory, which is submitted to the Office of Information Technology Services (ITS), and conduct risk assessments using the NIST AI Risk Management Framework. Strict privacy and data security guidelines are also included in the policy.

City Investments and Capacity Development

Berkeley has already committed significant internal resources toward digital transformation. The City's Information Technology Department has expanded cloud capabilities and data infrastructure. The Office

⁵⁵ "Consumer Protections for Artificial Intelligence." Consumer Protections for Artificial Intelligence | Colorado General Assembly, May 8, 2024. <https://leg.colorado.gov/bills/sb24-205>.

⁵⁶ City of Chicago. *AI Principles*. Department of Innovation and Technology. Accessed June 30, 2025. <https://www.chicago.gov/city/en/sites/chitech/home/roadmap-for-AI/ai-principles.html>.

of Economic Development has initiated pilot partnerships with local tech firms and universities.⁵⁷ Departments, including Planning, Finance, and Public Works, are exploring data-driven tools to increase responsiveness and optimize staffing.

Summary: Berkeley's Path to Responsible AI Governance

The proposed guidelines of the Berkeley Rule: Put Residents First; Modernize City Services; Empower Community; Standardize Operations; Ensure Transparency and Accountability; Certify Ethical Use; Protect and Prepare Our Workforce; Defend Civil Liberties; Social Advancement and Accessibility; and Catalyze Civic Wealth provide an aspirational framework for the responsible use of AI which prioritizes civil liberties, public trust and benefit with substantial human oversight. Such policies should anchor AI oversight in procurement discretion, ethical review, and internal use controls.

REVIEW OF EXISTING PLANS, PROGRAMS, POLICIES, AND LAWS

City of Berkeley Strategic and Policy Commitments

The City of Berkeley has adopted numerous strategic plans, ordinances, and administrative regulations that implicitly support but do not yet explicitly govern the responsible deployment of artificial intelligence (AI) within municipal operations. While these policies reflect strong commitments to transparency, social justice, privacy, and technological innovation, none currently provide detailed standards for the procurement, oversight, or ethical evaluation of AI systems.

Context at the State and Federal Level

California Consumer Privacy Act (CCPA): Provides baseline privacy rights for California residents, including the right to access, delete, and opt out of the sale of personal data. While not designed for municipal use cases, its principles inform best practices for consent, transparency, and data minimization in AI deployment.⁵⁸

Automated Decision Systems Accountability Act (AB 2930): The **Automated Decision Systems Accountability Act (AB 2930)** is proposed California legislation that would require government agencies to assess the risks and impacts of automated decision systems (ADS) used in public services. It mandates

⁵⁷ U.S. Economic Development Administration, *Regional Technology and Innovation Hubs Program*, accessed June 25, 2025, <https://www.eda.gov/funding/programs/regional-technology-and-innovation-hubs>.

⁵⁸ California Department of Justice, *California Consumer Privacy Act (CCPA) 5*, <https://oag.ca.gov/privacy/ccpa>.

transparency, equity analysis, and documentation of how such systems affect individuals, particularly in high-stakes areas like housing, healthcare, education, and public safety. Agencies would be required to conduct impact assessments, mitigate potential harms, and disclose the use and function of these systems to the public.

In alignment with AB 2930, the City of Berkeley could proactively collaborate with the **California Office of Data and Innovation (ODI)** to ensure that local AI deployments meet the highest standards of transparency, equity, and ethical oversight. This partnership would support Berkeley's efforts to pilot responsible AI practices, contribute to statewide standards, and share best practices for municipal implementation of ADS accountability frameworks.

ENVIRONMENTAL SUSTAINABILITY

The proposed resolution could advance Berkeley's environmental sustainability and climate resilience goals by guiding the ethical and strategic use of artificial intelligence (AI) within city operations. AI can support emission reductions, infrastructure efficiency, and climate adaptation through applications like smart energy management, predictive climate analytics, and digital permitting for green infrastructure. The framework ensures that these tools are deployed in alignment with the City's Climate Action Plan and Resilience Strategy, while centering on environmental justice and equitable access to environmental data.

Recognizing AI's substantial resource demands, particularly from energy-intensive model training and data processing, the resolution includes measures to assess and mitigate the environmental footprint of high-computation systems. Vendors should disclose projected energy usage and emissions, and the City should favor cloud-native, carbon-conscious, and open-source solutions that minimize resource consumption. These safeguards ensure that Berkeley's use of AI enhances, rather than undermines, its long-term environmental and equity commitments.

PROJECTED FISCAL OUTCOMES WITH AI

The fiscal impacts of implementing the proposed Artificial Intelligence (AI) governance resolution are expected to be moderate and manageable within existing departmental budgets, particularly in the early phases. The proposal is designed to leverage existing staff capacities, align with current procurement and digital modernization practices, and utilize external certification infrastructure rather than creating new regulatory bodies or bespoke municipal frameworks.

Artificial Intelligence should be considered for its potential to help staff workload and operational costs by automating routine tasks and streamlining service delivery. Establishing ethical safeguards, workforce protections, and oversight mechanisms may demand upfront time and coordination, but doing so is essential to ensure that AI deployment aligns with City values and avoids unintended harm.

Short-term costs primarily involve staff time for compliance integration, legal review, and development of an AI Use Registry, activities that could be absorbed through existing resources or targeted grants. Over the mid-term, minor capacity adjustments may support equity reviews and staff training.

Long-term benefits are anticipated through improved operational efficiency, enhanced revenue recovery, reduced legal exposure, and increased access to external innovation funding.

CONTACT PERSON

Councilmember Ben Bartlett	bbartlett@berkeleyca.gov
Councilmember Igor Tregub	itregub@berkeleyca.gov
James Chang	jchang@berkeleyca.gov
Olga Bolotina	obolotina@berkeleyca.gov
Artemisia Spencer Mace	aspencermace@berkeleyca.gov
Shamaila Ahmed	sahmed@berkeleyca.gov
Jayla Johnson	(510) 981-7130
Mia Surratt	(510) 981-7130
Nicole Liu	(510) 981-7131
Thelonious Spencer Mace	(510) 981-7135
Alejandro Ramos	(510) 981-7140
Samuel Chen	(510) 981-7140

Special Thanks to:

Kathryn Camp, Public Policy Consultant

kathrynmcamp@gmail.com

James A. Wolff, Chair, Emerging Technologies Law Group Warsaw Burstein, LLP

jwolff@wbny.com

De Kai, AI & AI Ethics Professor

hi@dek.ai

Peter Hirschberg, Director, City Science Lab San Francisco

peter@citysciencesf.org

ATTACHMENTS AND MATERIALS

1. Resolution

ATTACHMENT 1

Resolution establishing “The Berkeley Rule” artificial intelligence framework and vision to maximize public benefit through innovation and accountability.

WHEREAS, Artificial Intelligence (AI) technologies offer transformative potential to improve the lives of residents, enhance municipal services and operational efficiency, and support data-driven policymaking within the City of Berkeley; and

WHEREAS, responsibly deployed AI could streamline service delivery, reduce bureaucratic friction, improve emergency response, and enhance infrastructure resilience, while ensuring alignment with Berkeley’s values of transparency, sustainability, and social advancement; and

WHEREAS, the City acts solely in its proprietary capacity as a municipal service provider and market participant, exercising discretion over its own procurement, deployment, operations, and service delivery, and does not regulate or impose obligations on private-sector or individual uses of AI outside of City contracts or activities; and

WHEREAS, the City recognizes the risks associated with AI, including algorithmic bias, privacy violations, and procedural opacity, and is committed to mitigating these risks through ethical oversight, transparency, and public accountability; and

WHEREAS, the City of Berkeley hereby endorses "The Berkeley Rule," a values-based framework for the ethical adoption, procurement, deployment, and oversight of AI systems, ensuring alignment with Berkeley’s strategic goals and community values; and

WHEREAS, "The Berkeley Rule" incorporates the following principles to guide the use of Artificial Intelligence systems in municipal operations:

1. *Put Residents First*: Centering AI use on serving the health, safety, prosperity, and well-being of residents by improving access to essential services, reducing bureaucratic friction and eliminating unnecessary expense.
2. *Modernize City Services*: Leveraging AI with human oversight to upgrade service delivery. Ensure fairness and reliability by enhancing efficiency, responsiveness, and accessibility in City operations, including 311 services, permitting, and faster emergency response.

3. *Empower the Community*: Consider participatory governance models, such as an AI Oversight Advisory Board of diverse stakeholders to ensure ethical use, with public reporting, a Digital Ombudsman to support algorithmic review and redress, and an AI Sandbox to test new tools with community input, foster learning, and spark innovation and entrepreneurship.
4. *Ensure Transparency and Accountability*: AI Use Registries to provide residents accessible information about AI systems, their purpose, data use, and oversight mechanisms.
5. *Standardize Operations*: Operational consistency and safeguards, including centralized inventory management, streamlined procurement procedures, uniform oversight protocols, and adherence to robust cybersecurity and compliance standards, and rapid communication, to ensure responsible, transparent, and equitable adoption of AI systems.
6. *Certify Ethical Use*: Collaborating with independent organizations to establish AI ethics certification programs, ensuring voluntary vendor compliance with principles of fairness, transparency, environmental responsibility, and the preservation of humanity.
7. *Protect and Prepare Our Workforce*: Ensure that City employees grow alongside technological change by requiring Workforce Impact Statements for AI systems, providing retraining opportunities, and empowering existing workers via AI adoption.
8. *Defend Civil Liberties*: Build community trust, protect privacy, due process, and individual freedom. Prohibit unchecked surveillance, ban martial force, ensure residents control their data, encourage privacy impact assessments, and guarantee that all enforcement-related AI includes human oversight and the right to rapid appeal.
9. *Social Advancement and Accessibility*: Ensure AI expands opportunity and representation, reflecting the needs of all residents. Encourage equity risk evaluations for high-impact systems, pro-actively address harms and disparities, and design accessible tools through inclusive, community-led processes.
10. *Catalyze Civic Wealth*: Harness AI to optimize non-regressive revenue streams and spark new civic economies through municipal entrepreneurship. Generate public wealth for community reinvestment and deliver material benefits to all residents.

NOW, THEREFORE, BE IT RESOLVED, that the City of Berkeley affirms its values of ethical, aligned, and transparent AI use by endorsing "The Berkeley Rule," a municipal AI Framework for consideration that prioritizes commitment to innovation in service of liberty, dignity, and the public good.

