A California Strategy to Leverage Artificial Intelligence to Enhance Public Service Delivery in Local Government and Manage Risks

June 11, 2026

Jake Brymner
Policy Fellow
The Institute for California AI Policy (ICAP)
Silicon Valley Leadership Group (SVLG)

ABOUT THE AUTHOR

Jake Brymner is a Policy Fellow at the Institute for California AI Policy (ICAP) and a Capstone Fellow with the UC Berkeley Center for Civility & Democratic Engagement. Jake is the Founder and Principal of Jake Brymner Policy Solutions, where he supports public agencies and non-profit organizations in analyzing policy problems and developing evidence-driven responses. Jake holds a Master in Public Affairs and Graduate Certificate in Technology Policy from the Goldman School of Public Policy at the University of California, Berkeley.

ACKNOWLEDGEMENTS

This report was made possible thanks to several individuals and organizations. The author is grateful to Ziyang David Fan, Executive Director of the Institute for California AI Policy (ICAP), for his guidance and support throughout this project. Numerous individuals in both professionals and academic roles also offered instructive feedback during the course of this project, including public servants and members of the Gov/AI Coalition. Their candor provided insight that could not be captured through reviewing any amount of official publications, comparable legislation, or AI use inventories.

The author conducted this study as part of the program of professional education at the Goldman School of Public Policy, University of California at Berkeley. The judgments and conclusions are solely those of the author, and are not necessarily endorsed by the Goldman School of Public Policy, by the University of California or by any other agency.

EXECUTIVE SUMMARY

Many California local government agencies lack the capacity and/or systems to procure and deploy AI tools and manage their potential risks; as well as take full advantage of the benefits. As a result, there are missed opportunities to enhance public sector service delivery while existing uses of AI may not be transparent or assessed for potential bias and performance standards.

Local agencies in California are already utilizing AI tools in a number of ways, whether helping improve bus travel times, supporting caseworkers advising clients, or predicting residents at risk of becoming homeless. Understanding where agencies encounter friction points in procuring or deploying AI tools requires consideration of several elements in how uses differ: the problem being addressed with a tool, the dataset(s) involved, the product itself, and how it interacts with people and processes.

Review of relevant literature, a series of stakeholder interviews, and initial survey findings raise several challenges and potential solutions as public agencies navigate these new technological capabilities. Public agencies are inherently risk averse as they maintain delicate operations and address immediate resident needs. However, innovation with AI tools creates yet another set of challenges to navigate. Many agencies grapple with a lack of staff understanding of AI tools, as well as underlying data. Public agency workforces are also anxious about how AI adoption will impact staff, while management is cautious about triggering collective bargaining requirements around job duties. Procurement processes have not caught up with rapidly evolving AI uses, often failing to account for information on model performance or data sovereignty issues. Agency staff find that they often do not have all the information about a tool that would be helpful to understanding its impacts. Vendors experience challenges as well, finding procurement processes inconsistent and time-intensive before often being stuck in pilot phases. Deploying AI tools also creates issues for local agencies in costs, safety, and compliance for data storage.

California has taken early steps, but has yet to develop a comprehensive approach to support innovation and mitigate risks among local governments utilizing AI tools. Despite executive and legislative efforts, California's state-level policies have not fully addressed local adoption of AI tools. However, California does have some relevant foundations in its existing data services provided by the California Department of Technology (CDT) and recent actions to require certification of technologies procured by state agencies. States as distinct as Texas and New York have adopted state-level policies around adoption of AI tools by local governments.

Local agencies have already been crafting their own approaches in this vacuum. Several jurisdictions have approved local policies on AI use, while others have gone as far as to publish inventories of how they utilize AI tools. The City of San Jose has convened a national coalition, the GovAI Coalition, to facilitate sharing resources and experiences among agencies to inform future efforts. Larger jurisdictions have also begun testing means for learning about and testing capabilities of AI tools or contracting effectively. Beyond procurement, agencies are also implementing internal coordination and training to move toward adoption in alignment and understanding of tools' capabilities and risks. Local agencies starting now to define policies and strategic plans around AI adoption will benefit significantly from these peers and their examples.

However, local innovation can only go so far. Interview participants were clear on the impediments they confront and where state-level action can help enable constructive adoption and mitigate risks. California policymakers should advance the below package of policy and administrative actions to align state and local efforts toward leveraging technology for greater public impact of local government services:

Local policy & administrative actions

- Designate staff to serve as "AI Policy leads" and help coordinate interdepartmental efforts
- Create an internal training program to address the needs of both general staff that might interact with general purpose AI tools and "change catalysts" from across departments that are interested in helping identify specific instances where AI tools might be helpful
- Conduct an audit of current uses of AI tools and datasets upon which they rely or might be utilized in the future

- Draft and publish internal guidelines on how the agency will utilize AI tools
- Publish a public-facing inventory of AI uses
- Review and update procurement guidelines to ensure they reflect questions and issues specific to AI tools when appropriate
- Engage peer jurisdictions to understand where they have found constructive uses of AI and facilitate interagency knowledge transfer on practices like piggy-back contracting or sandbox testing

State policy & administrative actions

- Require that CDT issue guidelines on information that public agencies must receive when procuring AI tools from an external vendor
- Provide funding to CDT to expand training efforts and initiate communities of practice that are oriented toward local agencies
- Provide funding to CDT to facilitate a “sandbox” program which local agencies can test AI tools from vendors that volunteer to make their products available to use with synthetic datasets
- Initiate a pilot program in which CDT offers data storage and compute necessary for AI tools available to local public agencies
- Update data protection laws for data held by the state on behalf of local agencies

SECTION 1

Understanding AI in Local Government

Artificial intelligence (AI) has been in use across the public sector for years in various forms, whether as "predictive analytics," "risk assessments" or other iterations of "big data." However, technical developments through use of transformers in large language models (LLMs) and other forms of advanced machine learning have expanded the types of AI tools available to public agencies and their potential use cases. An OECD report found that 70% of governments had already deployed AI tools in a range of capacities, while the United States was host to the most use cases to be studied (OECD, 2025). Expanded uses of AI tools are emerging across all levels of governments.

As of April 2026, the "GovAI Coalition," facilitated by the City of San Jose, now includes over 750 state or local government agencies leveraging AI in their operations (GovAI Coalition, 2026). Other local governments in California are actively exploring how tools might be adopted; a survey conducted by and reported in Industry Insider found a large number of California cities prioritizing exploration of technology, and specifically AI, as a means of improving customer experience (Industry Insider, 2025). The extent to which both public sector consumers and vendors of AI tools foresee such significant uptake of AI tools in the years ahead raises the question: how are such technologies being deployed now and in what kinds of use cases?

Understanding the applicability and value prospect of AI tools in helping address public problems requires considering the efficacy of the tool itself, but also the context of its potential use case. Several elements to a local government use case facilitate assessing suitability, potential impact, and risk involved in deploying an AI tool. This section presents a taxonomy of the various kinds of AI used in local government settings and three example use cases from California.

A Taxonomy of AI in Local Government

Understanding how local government agencies deploy AI tools requires considering the problem being addressed, the dataset and product involved, and the interaction of people with the tool within the resulting workflow process. To better inform efforts to enable strategic adoption of tools and risk mitigation, policymakers must connect potential interventions to how these underlying dynamics manifest in procurement and deployment of AI tools. Policies will best advance sophisticated decisionmaking around AI tools when they are designed to help local agency staff navigate these factors within the practical phases for incorporating new technologies into operations.

The nature of the public problems that local agencies are confronting and in which they explore deploying an AI tool to address is crucial to understanding a use case. The stochastic nature of AI technologies enables their ingesting massive amounts of data to generate various kinds of outputs, beyond the capabilities of human staff. This makes many AI tools helpful for solving what might be considered "prediction problems" in which the new question to be answered is very likely to be similar to those that came previously (Kleinberg et al., 2015). This same capability to ingest data and provide near-immediate outputs given specified decision rules make AI tools helpful in addressing "efficiency problems" where highly routine, structured workflows are time-intensive for staff. Local government agencies also often face other kinds of "wicked problems," for which there are no obvious solutions or there is contestation around the desirable public response (Rittel & Webber, 1973). Though AI tools can assist staff in responding to these and public problems, they would currently tend to do so within sub-tasks related to either prediction or efficiency problems.

"Garbage in, garbage out" describes not only municipal waste management responsibilities but also the importance of underlying datasets to the outputs of AI tools. Public agencies must consider the sensitivity of the data used to train or being processed by a model, particularly whether it contains personally identifiable information (PII). AI developers often utilize "reinforcement learning" processes that seek to improve a model's future performance based on its current outputs, which can make an agency's outputs another ingested data

point. Access to the underlying "training" data to construct an AI tool, the data input into the model, and output data are thus all datasets that public agencies must consider in regards to storage and access, raising questions about data sovereignty and governance. Beyond the nature of data being collected, agencies must grapple with the collection process itself. Data might be collected in a manner that is inherently unreliable or subjective. Bias within a dataset is inevitable, as they reflect the biases in society more broadly and with data collection processes. More significant to how public agencies utilize a dataset is the extent to which its bias has been evaluated, interrogated, and understood. Understanding the structure, reliability, sensitivity, and bias of datasets training or utilized within a model is key to interpreting what comes "out."

The tools utilizing AI models have grown in their variability, as do their origin of construction and phase in the technology product lifecycle. Many forms of what we consider AI are now used in public sector settings, from classic machine learning prediction models to large language models like those offered via ChatGPT or Claude. Models vary significantly in their interpretability and the extent to which one can clearly follow the process through which an output is derived; some models are "black boxes" that preclude following a causal chain from input to output, while others are quite easily understood through a limited set of input variables.

Public agencies of sufficient scale may also face a question of whether to "build or buy" an AI product, and if the latter from which vendor. However, open-source AI tools present a third option that may become increasingly available as public agency capabilities and experience with AI tools grow. Even when exploring an AI tool, a public agency might begin to consider where it fits within a project lifecycle. Some exploration of an AI tool may occur within a "bake-off" context where a tool is actually deployed but not yet adopted. Agencies might then pilot or phase-in an AI tool before deploying it more widely or at-scale. After reaching scale, agencies should consider the maturity of the AI tool and whether it might become "stale" (i.e., degraded, eroded, or less accurate) or its outputs "drift" (i.e., different outcomes from what was previously expected) if not maintained. These issues around performance over time invite questions about how agencies would refine or modify models based on these issues, what one might consider "reincarnation" in the product's lifecycle. Policymakers must consider the nature

of the AI tool as a product across these dimensions in evaluating structures and supports for local agencies.

Despite the importance of the AI tool itself, its deployment in a real world setting, particularly that of a public agency, should be considered as a socio-technical system that interacts with both people and process. In a public sector context, AI models can either primarily serve users internal to the organization (staff) or external (residents, businesses, etc.). Models will also impact the public at different scales, informing actions or decisions at an individual or community level. The processes around how people interact with the tool provides another set of factors to consider. An AI-based functionality might be an entirely new one or might seek to improve upon an existing one. The functionality itself takes different forms as it overlays on these interaction points. A common means of considering how AI tools affect workplaces is by understanding their impacts as "automative" or "augmentative" in how they change existing staff processes; however, both represent what might be considered changes in work that are more technical in nature and do not fundamentally alter the design of program delivery.

Drawing from how Heifetz and Linsky differentiate technical and adaptive challenges, a third potential revision to process might be those that are "adaptive" rather than automative or augmentative, resulting in fundamentally restructured delivery from introduction of a new AI resource (Heifetz & Linsky, 2017). Questions of delegation, authority, and oversight sit at the intersection of people and process when organizations adopt new AI tools. Answering such questions requires greater nuance than simply identifying whether a "human is in the loop," but instead how they are. In some instances, staff might review every decision or output in a tool, while in others they might monitor the overall performance of the system or model. How decisions are reviewed or appealed is another key design element of a workflow.

San Jose: Running the Bus On-Time

In San Jose, local transit agency (Valley Transit Authority, or VTA) staff identified a missed opportunity to enhance bus service by using an AI tool to optimize its "transit signal priority" (TSP) system that prioritizes stoplight changes for buses. This challenge can be viewed as a fairly clear prediction problem (when will a bus arrive at a given intersection) but also as an efficiency problem (how to make existing transportation policy to prioritize buses at signals work better).

TSP already provides VTA with significant data on an ongoing basis, as do tracking measures for the timeliness of the bus lines, without any need for new data collection.

In the status quo approach, VTA had already delegated decision-making on traffic signals to an automated TSP system, so the adoption of a new AI tool represented an effort to improve an existing automated function and not any significant change in staff involvement. Like the prior TSP system, the tool itself was provided to VTA by a vendor, LYT.speed, and initially tested as a pilot before it was eventually scaled. LYT's model itself was built through a machine learning process and utilizes real-time information about the location of each bus to inform the timing of lights along a bus corridor, rather than the physical signals at stoplights used by traditional TSPs.

So far, VTA has reported that the use of the LYT system has resulted in on-time bus performance improving by 18-20%, positive impacts at a community level, though more concentrated for riders, and also reduced fuel costs (LYT.speed, n.d.).

Los Angeles County: Streamlined Benefits Delivery

A pilot in Los Angeles County has tested how caseworkers at social services agencies might benefit from access to a Retrieval-Augmented Generation (RAG) model that helps search and summarize technical documents about public benefit programs. The problem in this scenario is one faced by the staff rather than the public, specifically how to find the most relevant information for a client in as little time as possible to be able to then serve the next client.

The challenge faced by caseworkers then is one that is both related to efficiency (quickly finding relevant information) and prediction (will the information be relevant) problems. The dataset in this case is a highly stable, well-defined corpus of documentation on public benefits programs, though there is also a likelihood that caseworkers may enter some degree of PII into prompts when querying the tool. County agencies were piloting the use of this tool as provided by a vendor (Nava Public Benefit Corporation) and subject to research by independent evaluators from Georgetown University and Cornell University. This use case targets individual level impacts, but does so in a way that influences how both staff and clients engage; this tool changes staff interactions with documentation as part of their workflows, which in turn impacts how clients are treated. However, because this tool is ultimately used at the discretion of an individual caseworker, staff retain their authority of the interaction with clients.

When evaluating this pilot through a random controlled trial, researchers reported that caseworkers attained a 30% increase in accuracy of responses to clients when working with the assistance of a chatbot relative to those that did not, while also reporting reduced administrative burdens for staff navigating information on public benefits (Chen et al., 2026).

Another pilot effort in Los Angeles County leverages an AI model in a different manner, utilizing a large dataset to inform a tool to predict which county residents may be at the greatest risk of falling into homelessness. Working with the California Policy Lab, County staff used an existing dataset from the “LAHub” that connects data from across 12 different county agencies to test various means of predicting which residents might become homeless.

This large dataset was an existing resource within the county that contained sensitive PII data about residents' interactions with various agencies, including the county hospital and public safety agencies. After some testing, researchers ultimately decided upon using a fairly opaque model with 68 variables that were found to be highly predictive of which residents eventually became homeless. Using this information, the County then launched a pilot program that proactively offered at-risk residents with financial support (of approximately $5,000) and casework support from a special unit of county social services staff. This approach differs markedly from the existing services provided by the County, which requires that an individual first experience homelessness before receiving assistance. The model provides an output about which individuals should be contacted for support, but ultimately county staff must lead the outreach process. (Blackwell et al., 2024).

Results from the pilot program are still emerging, but so far appear positive: receiving support through the Homelessness Prevention Unit was associated with 71% decrease in residents' use of street outreach or interim housing relative to those who were not receiving support but had otherwise appeared likely to experience homelessness (Blackwell et al., 2025).

These are only three examples of the many growing applications of AI in local government, but they help illustrate significant variability in the kinds of tools already available and the potential for their positive impacts. Of course, other applications of AI in the public sector have already been attempted and not met their goals, such as the pilot of the "Elby" chatbot in Long Beach that responded to resident inquiries with old or inaccurate information (Descant, 2025). Or, worse yet, some uses of AI tools have produced negative impacts in communities or exacerbated inequalities, such as the use of predictive policing tools in Los Angeles (Lau, 2020).

These cautionary tales are also important to how we understand the role of local and state policy in guiding the trajectory of AI tools within the public sector. However, without many published inventories of AI use in local government available or disclosures on their impact, the bulk of available information tends toward agencies and vendors that are eager to share the positive results from pilot uses of AI tools. The lack of available information on the use and performance of AI tools will be further explored throughout this report.

SECTION 2

The current state of AI in California's local public agencies

Beyond documenting the existing status of AI tools in local government settings, this analysis seeks to contextualize their use within broader scholarship around how AI technologies impact organizations, particularly those in the public sector, the means through which government agencies procure and deploy AI tools, and how the underlying nature of AI technologies and local public agencies intersect. In addition to reviewing relevant literature, this section presents the findings from engaging relevant stakeholders through a series of semi-structured interviews and the initial findings from a survey released in late March 2026 to solicit further feedback.

Growing array of new AI tools, but inherent challenges to local government adoption

Despite the ready availability of a supply for AI tools built for government, there are unique impediments to thoughtful public sector adoption, particularly at a local level. Across the public sector, procurement processes have a reputation as overly complex, task oriented, and favoring incumbent vendors with prior or existing government contracts that can withstand the lengthy journey to entering an agreement (Burger et al., 2025).

Beyond general procurement challenges, nascent or novel technologies, like AI, can present additional issues. According to the National League of Cities, while some forms of AI are already being used in local government, more recent iterations of generative AI are still relatively unfamiliar to many working in municipal agencies. Because many local government agencies tend to default to prior contract templates and means for setting vendor tasks during the procurement process, local governments may face significant challenges when adapting to solicitations for impactful contract relationships for delivery of high-quality AI products and services (National League of Cities, 2024). While the GovAI Coalition has published some resources to help facilitate procurement and encourage vendors to report on information on performance of their tools, these remain non-binding and voluntary; staff have indicated that vendors tend to see "model cards" on performance of their tools as entry points to local agencies rather than documents in which to report on important factors to model performance (GovAI Coalition, 2026).

The stakes of these procurement decisions, often viewed as technical exercises, are significant as they reflect policy decisions embedded in the performance metrics for which models are optimized, datasets upon which they are trained, and structure of user interfaces (Mulligan & Bamberger, 2019).

In addition to any difficulty navigating dynamics with vendors, the context of local government creates other challenges in procuring, deploying, or overseeing AI tools. The National League of Cities convened an advisory committee and partnered with Google to explore key opportunities and risks associated with AI

among its members. Many of the barriers cited were of a political nature (organizational culture changes, sustainability of political support), while others were technical (data literacy and readiness, data security) or resource-driven (funding, access to necessary infrastructure) (National League of Cities, 2024). Cities across the country, including those in California, are facing significant budget deficits stemming from increasing costs and reduced revenues from tax bases and federal funding programs (Goodman, 2025).

Local government is also the most heavily unionized sector of the American economy (U.S. Bureau of Labor Statistics, 2025). Collective bargaining agreements restrict how duties for roles within local government can be amended to reflect new responsibilities involving the creation, use, or maintenance of algorithmic tools, limiting the pace of adoption and innovation (Goldsmith & Yang, 2024). Given the strength of organized labor in local government, it is in the interest of agency managers to consider how to develop staff buy-in to facilitate adoption of and utilization of AI tools, rather than leave it as another contentious point for bargaining – as is already occurring in some jurisdictions.

Local governments are inherently the most accessible for constituent engagement and hold higher degrees of public trust than state or federal governments (Jones, 2025; Nabatchi & Blomgren, 2014). This trust is both an asset to public leaders seeking to innovate through use of new tools, but also presents risks for those that fail to disclose the use of AI tools or adequately test them before deployment. Local public agencies may be eager to explore new means to enhance their impact, but must do so within sector-specific fiscal, political, and technical constraints.

Beyond these challenges in local government, AI tools also create or exacerbate unique risks relative to other types of technologies. AI tools are only as effective as the data upon which they are built are reliable. Issues in the means of data collection or recording by government can leave AI tools susceptible to statistical errors, while other structural issues can create systemic biases in the outputs of AI tools that reproduce existing disparities between communities (Feng & Wu, 2019; Xiang & Raji, 2019).

These issues raise the stakes for how public sector agencies will determine who is accountable, or liable, for the outputs of AI tools. Software products are generally not subject to existing liability frameworks, which heightens the stakes for how procurement and contracting processes clarify the relationship between vendors and government consumers. Within organizations, and especially local government, AI tools also stand to alter the existing lines of accountability and discretion (e.g., between front-line staff and technology tools, or between staff and managers), while potentially changing interactions between agencies and the public (e.g., creating more accessible interactions with agency outputs).

California: Home to AI developers, charting a path for AI policies

California is home to most of the world's leading AI developers, but has been uneven in its approach to public sector adoption of technology tools. This issue is not unique to AI tools, but is notable given the extent to which California's elected leadership have articulated a desire to coordinate use of AI in government. As part of an Executive Order issued in September 2023, Governor Newsom directed state agencies to develop a strategy for how to integrate GenAI tools into operations (Office of Governor Newsom, 2023).

The Little Hoover Commission issued a December 2024 report focused on how California state government agencies should prepare for AI adoption, but the resulting policy recommendations would also stand to benefit both state and local public agencies. The Commission proposed construction of a California state compute center to provide processing, data, and support services to agencies, as well as establishing an AI Council that could develop training, guidelines, and procurement resources (Little Hoover Commission, 2024). Senate Bill 53 (Wiener) mentions a "CalCompute" center, but it has not been funded and has a more narrow academic focus than what was proposed by the Commission.

Governor Newsom's executive actions and ongoing legislative efforts represent meaningful early steps, though a comprehensive state policy framework for AI in local government is still taking shape. In April 2025, Governor Newsom and local Los Angeles City and County leaders announced a public-private partnership to utilize a new AI tool to expedite building permits for construction following the Eaton and Palisades Fires. With philanthropic funding support and policymaker commitments to streamline review processes to help, local leaders identified a company (Archistar) that designed a pilot tool for assessing compliance with relevant zoning guidelines (Office of Governor Newsom, 2025). However, this more so represented the state helping facilitate or promote a pilot tool made available to local agencies, rather than any substantive involvement in its eventual deployment or financing.

Later in December 2025, Governor Newsom appointed an Innovation Council with a "Modernizing Government Services Delivery" subcommittee, but the effort

remains staffed within an existing state agency and it is unclear to what, if any, extent intended focus will include local government operations (Office of Governor Gavin Newsom, 2025). Similarly, the Office of Data & Innovation (ODI) facilitates a community of practice for innovation in state agency operations, but this venue is not open to participants from local government (Office of Data & Innovation, 2026).

Most recently, Governor Newsom issued an Executive Order in March 2026 on "Trusted AI Procurement" after the Trump administration declared Anthropic as a "supply chain risk" following a public dispute over the scope of use of its products by the Department of Defense (Albergotti, 2026; Office of Governor Newsom, 2026). Governor Newsom's order decoupled California's procurement process from the federal designation of supply chain risks by allowing for the state's Chief Information Security Officer to determine that such products can still be safely utilized by state agencies. Additionally, the order creates a new certification in the state contracting process in which firms must explain how their technologies are governed to reduce risk of "harmful bias," and avoid utilizing illegal content (such as child sex abuse material) or violating civil rights. Though these recent actions present helpful precedents for state involvement, it remains to be seen on how they can substantively advance how California interrogates issues in AI tool performance beyond minimum existing legal requirements or high-level statements of intent.

Data storage and compute capacity are other areas where California can adapt its state services to support local agencies. Though California has state-managed data storage infrastructure, it has not been oriented around needs specific to AI deployment or structured to align the fiscal incentives of local agencies. Nearly 4 of 10 state chief information officers (38%) reported in a 2025 national survey that their states offer data hosting to local agencies (National Association of State Chief Information Officers, 2025). Since 2014, California has been among the states that offer such services through the California Department of Technology (CDT), providing a range of options for state and local public agencies, including services for data storage. Despite larger private providers (e.g., Amazon Web Services, Azure) offering more competitive pricing, CDT has seen dramatic growth in subscriptions, from 137 in 2020 to 5,930 by 2025, a 43x rate of increase (CDT, 2021 & 2025). Much of this is likely attributable to the broader growth in

public sector cloud computing during this period, particularly following the pandemic-era push for increased access to digital services (Geraghty, 2021).

However, many jurisdictions continue to rely upon private sector data infrastructure. Public agencies often lack the personnel necessary to solely utilize data storage without staffing services offered by Managed Service Providers (MSPs). In a 2023 survey of city and county IT leaders, 36% indicated that their agencies retained MSPs (CompTIA, 2023). The pricing model for CDT's services is differentiated based on the volume and activity level of data being stored, without consideration for the nature of data. CDT's data storage costs are significantly higher than those offered by the private sector, no matter the volume or activity level (Legislative Analyst's Office, 2022). Currently, CDT offers the same pricing structure for agencies regardless of the potential risks associated with such data being hosted elsewhere. This approach incentivizes resource-constrained public agencies to approach data containing PII from a cost containment standpoint, rather than considering the security and privacy of such data as a public good.

The enactment of formal legislation governing AI and public sector use cases remains an evolving area for California. Broad-based proposals to regulate automated decision systems (ADS) in "high-stakes" decisions have prompted substantial legislative debate about scope, implementation feasibility, and potential impacts on innovation, and have not advanced through the full legislative process. More targeted legislative efforts to address AI and ADS tools in state government have had mixed results. A 2023 law requiring state agencies to report use of "high-risk" AI systems resulted in none being inventoried, a finding that, alongside a fiscal analysis citing costs in the "tens of millions" for proposed ADS risk mitigation requirements on "consequential decisions," illustrates the implementation challenges of broad-based AI reporting mandates (Johnson, 2025). Other legislation (Senate Bill 892, 2024) sought to regulate how state agencies procure AI or ADS tools, but did not include local agencies and was vetoed by the Governor.

California has previously acted to regulate how local agencies disclose the use of vendor-provided technology systems. Senate Bill 272 (Hertzberg, 2015) required that local public agencies maintain and publish an inventory of the "enterprise

systems," defined as "software or applications or computer systems," that includes information on its purpose, the vendor involved, and associated data collection and use practices. State law provided some exemptions to the nature of systems subject to these disclosures, explicitly exempting emergency services or infrastructure control systems for example, and offered a lower threshold on required disclosures in instances when the "public interest" is better served by not releasing additional information. These exemptions came about following opposition to the legislation from the associations representing county, municipal, and other local agencies, which expressed concerns regarding cybersecurity risks associated with disclosing information about systems used by emergency services and critical infrastructure (Assembly Committee on Local Government, 2015).

Despite this broad exemption, California has also enacted state law specific to how AI is utilized in policing. In 2025, California enacted Senate Bill 524 (Arreguin), which required that public safety agencies adopt policies to disclose the use of AI in the drafting of any official reports and data on the output provided by the AI tool and an audit trail on how it was revised, which was too opposed by agencies. Taken together, these two laws reflect the limited role of California legislation on local government uses of technology systems and AI tools thus far.

Nationwide, state governments across the U.S., and its political spectrum, are beginning to enact legislation to orient public agencies around AI adoption. In December 2024, New York enacted the "LOADING Act" (or, the "Legislative Oversight of Automated Decision-making in Government" Act) to require that state agencies publish inventories on AI uses including the name of the model's developer, its function, and specific use case, as well as require that agencies conduct impact audits on any new AI tool. Last year, New York extended these requirements to all local jurisdictions within the state, though agencies are still within the initial implementation period and yet to confront the deadline for compliance (New York State Senate, 2025).

Texas also enacted two laws relevant to its statewide approach to AI governance in 2025. That May, Texas adopted House Bill 2818 to create an AI Division within its Department of Information Resources (DIR), to be led by a new Executive Director position. Several weeks later, Texas enacted Senate Bill 1964, which

established requirements not only for documentation of both state and local agency uses of AI, but a number of other activities to be conducted by the DIR Resources. Senate Bill 1964 requires the DIR to issue an AI code of ethics and guidelines to be adopted by all public agencies.

DIR is also required to provide training and other educational materials for public agencies, as well as support an “Advisory Board of Public Sector AI,” appointed by the Governor. The Advisory Board was charged with considering potential uses to reduce administrative burdens while identifying regulations impeding adopti om or opportunities to streamline procurement. By March 2026, the Advisory Board had met twice and was reviewing a “sandbox” pilot program to aid in exploring AI tools (Texas Department of Information Resources, 2026).

To date, most local jurisdictions in California have not provided public information about uses of AI tools in their operations. An assessment of California’s county agencies and the 25 largest municipal jurisdictions found that only 12 counties and six cities had published policies on the use of AI. Among those agencies, only two counties, one city, and one consolidated city-county had published an inventory on their uses of AI within agency operations. While smaller municipal jurisdictions were not studied, feedback from stakeholders suggests that smaller jurisdictions are less likely to have any official policies adopted due to resource and capacity constraints.

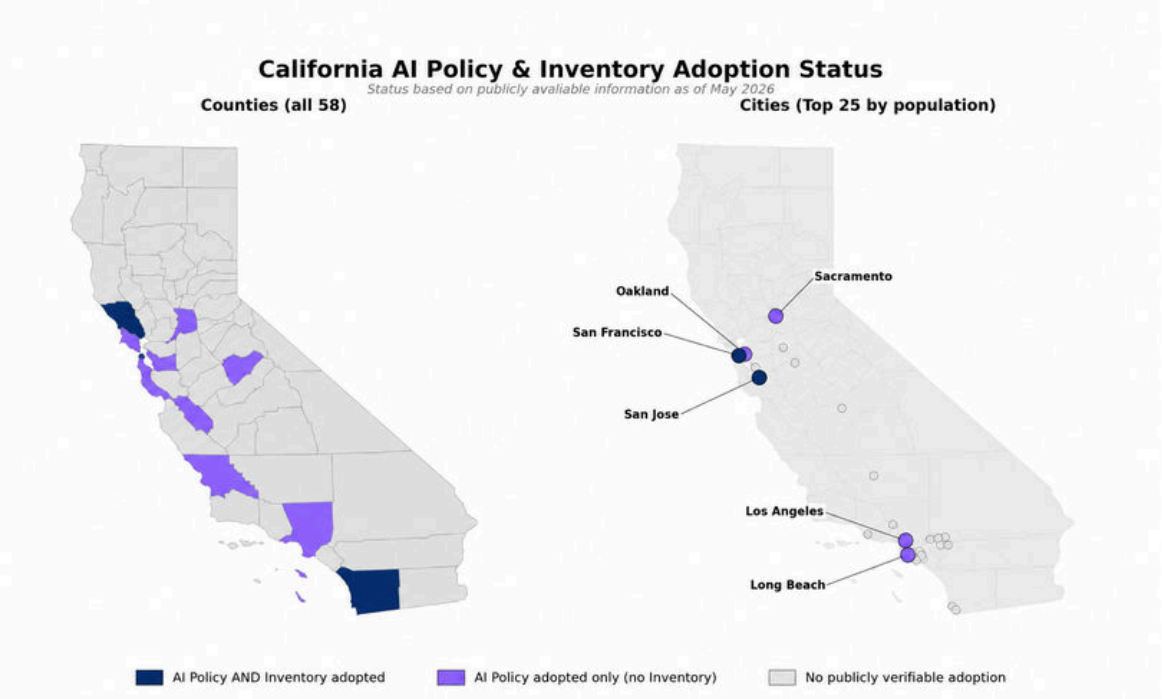

The lack of public disclosures by local jurisdictions on how they are utilizing AI tools complicates research into their potential impacts. Agencies that are communicating how they are utilizing AI tools or the frameworks guiding their decisionmaking around adoption are not necessarily typical of California cities or counties. These agencies are larger in size and serve more affluent communities relative to California that have more interaction with the technology industry, potentially influencing the level of expectation that residents might have in their local government's disclosure on uses of AI tools. As noted in Section 1, there is also a bias towards reporting on uses of AI tools that are successful, rather than those that are explored, tested, or piloted but do not move forward or remain in deployment. Despite this constraint, there is still value in exploring how California agencies today are already deploying AI and applying the framework presented in this paper to understand their implications.

Results & findings from stakeholder interviews & feedback form

Between February and April, the author conducted 15 semi-structured interviews with practitioners and experts with knowledge about how AI tools are being deployed by local governments. Interview participants were primarily staff in local government agencies who are leading the development and implementation of efforts to facilitate or guide the adoption of AI tools, several others were vendors of AI tools, provide technical support to public agencies, or research how agencies utilize AI. To encourage candid participation in these interviews, participants were provided with the opportunity to offer perspectives without direct attribution in this or other publications.

In addition to these interviews, the author also constructed a short survey, administered via Qualtrics, to solicit insights from government practitioners and vendors around key research questions and perceptions about potential policy responses. The GovAI Coalition shared information about this survey with its membership in a newsletter distributed in late March. The survey remains open at the time of drafting this report and will inform other publications, but initial responses provide additional perspectives that reinforce key themes from the interviews conducted during this research. Taken together, these interviews and initial survey responses provide more timely assessments than existing literature given the quickly evolving nature of public-sector AI tools and uses while also offering examples of promising practices and policy opportunities.

Interview participants from both sides of the procurement process identified parallel challenges. From the agency standpoint, participants expressed concern that procurement is often a vendor-led process, with one going so far as to say that agencies were “inundated with vendor pitches.” The volume of requests worried one participant about the longevity of vendors, which weighed on decisions around whether to contract. Despite the strong interest from vendors, agencies found that they often encountered barriers in receiving sufficiently detailed information on tools. When vendors do utilize voluntary informational tools, such as model cards, fact sheets, or platforms there are ranging approaches to disclosure; some vendors approach such documents primarily as business development tools while others provide more technical details. Generally there

was agreement that the procurement process was cumbersome and agency documentation had not evolved to solicit information that was specific to the nature of AI tools.

Vendors too expressed frustration with procurement as overly long and often not transparent about the agency's intention around the potential duration of a contract. Smaller developers expressed frustration with inconsistent types of information sought across agencies, which they claimed larger incumbent technology companies are better able to navigate with resources sufficient to retain more business development staff. A lack of agency staff expertise also meant that vendors saw agencies as highly reliant on specific individuals to navigate procurement. One participant believed that California had not yet acted on legislation around procurement due to high-visibility legislation seeking to regulate all automated decision systems rather than focus on public sector uses. Initial survey responses aligned with these insights, with respondents indicating that securing sufficient information from vendors was a challenge and disclosures were inconsistent.

There were also several insights into how procurement, or development, of AI tools was already being improved upon as agencies moved forward. Several agency staff raised models that were being tested within their jurisdiction, including several that stood to benefit their peers in other communities. One participant shared that their agency was involved in a contract under development using a cooperative purchasing approach, in which multiple agencies would enter into the same agreement for services. This multi-jurisdiction contract would provide consistent rates to all public-sector parties while also allowing for a "piggy-back" option for other agencies to join later at the same rate.

Several participants were also very optimistic about how the public sector might build its own AI tools, particularly given the rapid improvement of resources like Claude Code. The diffusion of application development resources was seen as an enabling factor for more public agencies to elect to "build" rather than buy tools, particularly if there was a highly repeated pattern of workflow across multiple departments. This improved feasibility for in-house development is somewhat curtailed by challenges with the cost of tokens for use of AI models that staff are

using to assist development work. One participant summarized this current scenario as agencies needing “permission, not capacity” to create more tools. Because local governments do not compete as industry rivals, but often seek to empower neighbors to address shared governance challenges spilling over borders, there was also an interest in sharing open-source government-built tools across jurisdictions.

Other promising means of testing new AI products were “bake-off” tests in which agencies utilize multiple tools during a pilot or testing process before making a selection. Similarly, “sandbox” demonstrations enable agencies to utilize a tool in a more controlled environment to assess whether it would suit broader operations.

Interview participants also saw a need to address internal issues within public agencies to better prepare for AI adoption. Nearly all participants identified a gap in AI literacy among staff, including senior leadership often in a decision-making role about operational changes. Even before literacy around AI specifically, some departments are struggling with digitization and are still highly paper-based. Similarly, a gap in understanding of data collection and management practices were cited as another stumbling block for many agencies.

Agency participants indicated that while many staff are eager to learn more about AI tools, others are anxious or even fearful. Concerns about automation and displacement are not uncommon, while other staff worry that they might be perceived as less useful if they do not use AI tools available. At an organization level, agencies are strategically approaching how they share AI tools and encourage their use among staff, carefully connecting guidelines on new AI tools to existing job duties. Some interpretations of state employment law suggest that any changes to staff workflows may constitute a change in job duties that triggers a “meet and confer” requirement between the agency and its labor unions. Agency participants suggested that staff training is helpful, but best received when offered by an in-house subject matter expert at the agency rather than an external provider.

Individual level gaps in knowledge or trust are exacerbated by organizational dynamics within public agencies. Nearly all participants spoke to some degree

about the pace of innovation within public organizations, whether characterized as a warranted need for stable operations or as risk aversion and a "permission-based culture." One participant compared how agencies respond to innovation as an "immune response" from the system to what is perceived as potential risk. Another technical assistance provider shared that many agencies are developing AI policies not to identify areas for adoption or innovation but as a means of controlling for risk and compliance, often from conversations initiated by City Attorneys or related roles attempting to manage potential liabilities for their agencies.

The positionality of where guidance and policies emerge matters to how they are felt across a jurisdiction. Generally, most agencies indicated that they are seeking to set policies and offer resources to other departments. One participant described their unit as having only "soft power" relative to other departments, seeking to encourage behavior rather than being able to issue mandates given a lack of enforcement levers. In another jurisdiction, a "Data Governance Committee" has become a de facto coordinating body around AI policies to avoid redundancy as it already retains the most relevant staff; this group has found a direct connection to the City Manager as important to how their work is shared across departments.

Interview participants were vocal about how data issues are impacting adoption of AI tools in a number of dimensions. Many agencies find that there is a lack of standardization in data quality practices across their departments, including a limited inventory of available datasets and issues when attempting to share data even within the jurisdiction. What happens with data that leaves the agency is another matter. One technical assistance provider shared that many agencies are not sufficiently inquiring about what model developers are doing with data inputted into a model and shared externally, often not asking about this during the procurement process.

One agency participant flagged how greater magnitudes of data generation and storage are creating new cost and compliance issues. While image processing tools rely on large amounts of video footage, they only utilize a fraction to inform decisions or action. However, current public records law is opaque in how agencies should treat this data, prompting retention and reporting issues while

also yielding significant costs. At this agency, even reducing storage time for one set of image data to 30 days resulted in $140,000 in annual savings. Local government staff also expressed a concern about the costs of compute required for AI tools and lack of means for public agencies to strategically contract to most efficiently meet their needs. Many public agencies are starting this work with a lack of data infrastructure and governance, with more potential cost and compliance issues on the horizon as potential uses of AI tools expand.

Agencies with dedicated staff and formal AI policies are intentional about how they are finding areas where tools might advance their impact. Most agency staff participants expressed that identifying problems or areas in workflows where AI might be helpful needed to start within the “business units” that are closest to an issue, rather than coming from other departments. Several participants described “change catalysts” within business units who were often looking for means of improving upon existing processes and seeking out information about AI tools that might be of use. Agencies with AI staff leads expressed that they do learn about potential use cases and tools when attending conferences or learning from peers, such as the network facilitated through the GovAI Coalition.

One more advanced agency had developed a framework through which they identify the decision-making unit and decision-making process for each area where new tools were being deployed in order to clarify staff roles and avoid the potential for delays from unforeseen veto points. Another agency had recently launched a new resident-facing program that enables community members to propose how new AI tools might help improve municipal services and then selects several projects to receive staff support in piloting. Many agencies are already developing internal structures to help facilitate exploration of innovation opportunities that might be enabled by AI.

Use of AI tools was consistently occurring at two levels. In one approach, as described above, a specific challenge is found and the overseeing department works with IT or AI policy leads to explore suitability of new tools. However, every agency participant also described a decentralized approach in which staff are given access to enterprise level LLM tools, like Microsoft Copilot or ChatGPT, provided with general guidance about how to use such tools, but not limited in how they might interact with them. Interview subjects often spoke about how

frontline staff needed to explore for themselves how general purpose GenAI tools might be helpful in their workflows, without overly restrictive guidelines dictating their interactions. Both these approaches were often couched in language about letting innovation occur on the margins before eventually seeking to scale it across the agency or with other staff. Agency participants also often spoke about understanding where staff were dedicated the most time or encountering the most frustration and exploring where AI tools might help overcome such challenges – not only as a means of making those particular tasks more efficient, but also to build greater support for partnerships within the jurisdiction and trust in the underlying tools.

Two interrelated challenges for local governments: AI tools market challenges (procurement) & changement management (deployment)

A review of the current state of AI deployments in California's local agencies, as well as stakeholder feedback via interviews and initial survey results, demonstrate that there are two interrelated tracks of problems faced by local agencies. Challenges in the procurement of AI tools by local governments stem from market-based problems from the unique nature of public sector consumers, while effective deployment of these tools often confront organizational change management problems. However, better procurement is facilitated by considering deployment, while successful deployment is contingent on an effective procurement process. Public sector agencies, and the vendors serving them, must ultimately consider both these dynamics in tandem. The stakeholder engagement process for this research confirmed and builds on existing literature documenting where agencies encounter friction points in both procurement and deployment across the four elements of the taxonomy of use cases presented in Section 1. These same friction points not only impact public agencies, but also the vendors offering AI tools as well as the public relying on services provided by local government.

Market problems for AI tools for local government use arise from both issues specific to these public agencies as well as the nature of emerging technologies. The highly regulated, and often convoluted, procurement process for the agency seeking to secure a product or service might be considered as creating increased transaction costs relative to other technology product markets. Relatedly, this presents a more opaque and time-intensive process for developers seeking to make products available to consumers, resulting in increased transaction costs for these vendors and creating a bias toward larger firms with existing resources, whether capital or dedicated teams oriented around procurement or public business development.

The lack of existing staff subject matter knowledge within many agencies and inconsistency in information available about AI tools also creates an information asymmetry in which government agency consumers lack the same extent of knowledge as firms. A culture of risk aversion also means that many public

agency staff and leadership have reduced incentive to explore adoption of AI tools. Together, these factors in the public sector market for AI tools create greater friction for local agencies in exploring adoption as a strategy for improving service delivery.

Table 1: Two sides of one coin: how challenges across the taxonomy of AI uses in public sector agencies arise in their procurement and deployment

	Procurement	Deployment
Problem	• Tools presented as solving problems they do not	• Tools used for a function for which they were not designed
Data	• Lack of clarity in training data, treatment of PII • Misalignment between training data and data utilized in practice	• Loss of data sovereignty • Lack of transparency on use of output data in model training • Costs of data storage
Product	• Inconsistent information available about tools • Process not oriented to solicit necessary information on tools	• Weak oversight of AI tool over time • Lack of transparency on use of the tool • Lack of interpretability of the model
People & Process	• Risk aversion • Lack of staff understanding of AI tools	• Resistance to changing workflows, number of veto points to changes • Lack of staff understanding of AI tools

Even after navigating procurement, local agencies face further hurdles relative to other organizations as they move forward in a change management process to integrate tools into existing operations. The limited staff understanding of tools, particularly taken with anxiety about how tools might displace labor, are barriers to develop workforce buy-in. Agencies must also consider collective bargaining requirements when assessing whether there are changes in job duties or employment conditions resulting from incorporation of AI tools.

A lack of understanding of the technology generally or the tools specifically can also result in an erosion of trust, both at a staff level and with the broader community served by that local agency. Managing this trust factor requires that agencies understand how data is collected, stored, and analyzed for bias, which local jurisdictions may find beyond their current capacities. Without an understanding of this data, agencies may not build processes to sufficiently monitor and remediate AI tools or the operations utilizing them. The nature of local government services, which are often authorized or funded by state or federal programs, also can restrict the options available for innovative uses of new tools.

Many of the change management issues stem from failure to properly plan for deployment during the procurement stage. Conversely, agencies that have not reinforced their internal ability and willingness to embrace new technology tools through a robust understanding of change management will not be well-suited to identify and select appropriate vendors. Local public sector leaders must recognize that procurement of new technology like an AI tool is a policy undertaking, not just a contractual transaction.

SECTION 3

Guiding constructive uses of AI tools in local public agencies

AI has already come to California's public services. Now, policymakers at both the local and state levels must ask how they intend to guide its growth and usage toward public impact and away from potential risks. Promising practices emerging within several jurisdictions provide other local governments with examples ready for replication. Local leadership and focus will prove most important to how AI deployment fares, so the proactive work by early adopters to devise strategies, frameworks, and policies should be encouraged by state leaders.

However, the challenges identified in this research suggest that such local action is necessary but not sufficient. While proponents of subsidiarity might suggest that state policymakers allow local agencies to determine which tools are appropriate to support service delivery, this paper concludes that several state-level policies can facilitate continued innovation while also mitigating potential risks, beyond what any individual local government can realize through its own action.

Advancing both innovation and safety, while also balancing multiple often competing priorities, requires that policymakers and others interested in leveraging technology for public impact must ask several interrelated questions:

- *How should policymakers weigh potential responses to AI in local government?*
- *Which policy levers can help enable responsible innovation in local government?*
- *What actions today can guide AI deployment toward better impacts tomorrow?*

The following section of this paper is organized around potential answers to these questions and their implications for California.

How should policymakers weigh potential responses to AI in local government?

In considering potential responses to address issues in the adoption of AI tools by local government agencies, California policymakers must hold several important criteria when evaluating likely impacts. As described in Section 1, AI tools are being deployed across a myriad of different types of uses that can complicate how any policy response fully considers the implications of each unique instance. However, the following criteria should guide policymakers in identifying means to address public agency capacity to adopt or deploy AI tools generally. Some of these criteria may be more significant in deliberations at a state or local level, given their relative proximity to address certain dimensions of the issues documented in this report, but each represents an essential factor to incorporate into a decision-making process on potential responses.

First, any policy should increase the capacity of public agencies to understand issues specific to adoption and deployment of AI tools and to monitor their performance over time. Not only must public agencies have sufficient capacity in order to effectively procure or adopt AI tools, but they must also have an understanding of the underlying technology and the specifics of the tools that vendors are providing. Policy responses should seek to both enhance public sector capacity and increase the information available to such agencies.

Progress on this criterion could be assessed by considering the number of staff at local agencies that are familiar with AI technology and issues relevant to gauging applicability of an AI tool to addressing a public problem or workflow, as well as their ability to ask questions of prospective vendors and to effectively secure relevant information about the tool's performance. The focus of policymakers should be empowering local governments to better navigate opportunities and challenges to service delivery from AI tools, so the effectiveness of any policy in building this capacity is the most important criterion in evaluating options.

Public agencies in California approach their work from highly varied starting points in regard to fiscal conditions and community needs, often creating an imbalance in which those serving the most vulnerable populations do so with

the fewest available resources. Technology adoption and oversight requires an up-front commitment of resources, even if only staff, that can further compound this inequity in resource access. Differences in geography, urbanicity, and political composition create an even wider distribution of potential starting points to consider for local public agencies, each potentially widening inequities that can further undermine effective adoption and oversight. Local agencies with greater resources (e.g., dedicated staff positions, in-house expertise, funding) are better enabled to more comprehensively consider how AI tools can enhance service delivery and undertake precautionary steps in adoption or deployment to mitigate risk.

Any policy should seek to reduce such gaps in capacity that exist across public agencies based on these contributing factors. At a local level, agency leaders should carefully consider how various communities, particularly the most historically underserved, will be disproportionately impacted by how an AI adoption policy is crafted. Measurement for this criterion can be assessed based on the number and diversity of California municipalities that become able, or are compelled, to undertake new activities, whether procurement, adoption, deployment, or disclosure, relative to the use of AI tools.

These inequities exist across agencies along a number of factors and are not unique to adoption of technology tools, stemming from more fundamental structural issues in public financing and California's demographics. Policymakers should design responses informed by an understanding of resource inequities across and within local jurisdictions and find opportunities where possible for an alternative to contribute toward reducing or offsetting such a gap, but refrain from viewing a policy as a panacea for much more systemic root causes.

While resource constraints may vary across jurisdictions they are prevalent, to some degree, in all of them. Reducing costs, whether in nominal dollars or staff time, is thus an essential element to any policy being successful and sustainable toward enhancing public agency capacity. Currently, the time required to study opportunities for how AI tools might enhance service delivery, solicit vendor applications and then select, deploy, and oversee the use of a tool can be prohibitive. There is also a significant degree of duplication that occurs across jurisdictions as each navigates its own procurement process across the same kinds of public problems and universes of potential vendors.

Vendors too bear costs when confronted with redundant requests for information from public agencies or when presented by process-driven procurement requirements. Procurement is a common source of friction in public sector agencies, but a uniquely acute impediment for AI tools relative to other government purchases given the specific types of information necessary to understand the potential impacts of AI tools. Similarly, though recruitment and retention of information technology staff has historically been challenging, the relatively novel nature of AI tools for public sector uses also means that there are fewer incumbent agency staff versed in AI-specific issues that matter for its evaluation and training.

Policies should seek to reduce these cost factors without creating incentives for agencies to bypass necessary due diligence or shirk oversight responsibilities. The ability of a policy to reduce costs might be measured by tracking any reduction in time required to navigate procurement and savings in staff time to develop internal policies, model performance tracking protocols, or other work involved to manage AI tools. Because resource constraints are often cited as a barrier to adoption and effective oversight, policymakers should carefully weigh how any response improves the efficiency of efforts to harness AI.

Policy responses to AI adoption by public agencies should also promote greater transparency and accountability, both for local governments deploying such tools and the vendors equipping them. These are interrelated priorities as public agencies can only communicate what they themselves know about a model, while vendors are more strongly incentivized to disclose information about their products if prospective clients are subject to mandatory reporting. Expectations for information disclosures can also prompt more regular evaluation of the uses and potential disparate impacts of AI tools, avoiding the erosion of public trust if problems emerge in program delivery and AI tools are later found to be involved.

Measuring impacts on transparency will differ between public agencies and vendors. For local governments, the number of agencies reporting on the use of AI tools, the number of elements upon which they are reporting, and how accessible this information becomes to the wider public are viable means of forecasting and tracking the efficacy of a policy. For vendors, transparency can be assessed by the level of information publicly available on the tools offered by

vendors and the number of vendors that provide documentation on the design, performance, and risk or bias in their models. Transparency is fundamental to government operations, especially at the local level and essential for determining the suitability of an AI tool in addressing a public problem, making this an important consideration for both state and local policymakers.

California is home to the largest AI development ecosystem in the world, and faces the challenge of developing governance that advances the needs of both its innovation economy and broader society. This challenge has been particularly true of efforts to establish state-level legislation on AI development, transparency, or its use in public sector operations. As noted in Section 2 of this report, Governor Newsom has issued several Executive Orders pertaining to AI and its use in state government, but has vetoed several bills passed in the state legislature by wide margins. Veto messages issued by the Governor typically articulate either costs to the state or that proposed restrictions would stifle innovation and ongoing state efforts.

Given that voluntary mechanisms for information disclosure and resource sharing already exist, the formal adoption of local or state policies is a critical next step toward providing greater clarity and predictability to agencies and vendors alike while also necessary to achieve other criteria for successful outcomes in expanding local government capacity. The feasibility of any policy to secure legislative enactment, budgetary appropriation, or inclusion in an administrative action is thus a significant factor in determining its potential impact in augmenting local government capacity relative to AI adoption.

Which policy levers can help enable responsible innovation in local government?

Both local and state policymakers hold several levers that can help guide how AI tools are brought to bear within local public agencies. Some levers exert greater force across certain dimensions of successful public sector capacity building, while failing to address others. Some levers may have significant long-term benefits, but relatively low initial impact. Though levers deployed locally will be felt most deeply within the jurisdiction at hand, jurisdictions like San Jose have already provided significant benefit to others across the country in the resources they have published based on their experience and engagement with other early adopters. What starts as one agency's local action to manage its own AI adoption and deployment can very well manifest into national best practices given the early stage of governance of these tools and dearth of formal regulations.

Conversely, though the scale of California means that state levers will have the furthest reach, most of the issues involved in adoption and deployment of AI tools are ultimately at the discretion of individuals within local agencies. Relying upon compliance with state laws has rarely proven the most effective strategy for enabling local innovation. Thus, this paper identifies a strategy for deploying policy levers in coordination, across local jurisdictions with state standards and resources available, in pursuit of a commonly held vision for how AI tools should be used for public impact.

Another means of considering where policy responses will exert the most leverage is assessing how they might influence either asymmetric information and "transactions costs" in the procurement of an AI tool or change management issues described in Section 2. Some potential policy levers may only help improve information available to local agency leaders as they consider whether a tool is suitable toward aiding in responding to a public problem. Others might help better enable agency staff to track and monitor the performance of an AI tool or communicate about its usage with interested stakeholders, but do little to improve what kind of information is available when agencies decide upon a particular AI tool. Several of the policy options presented

here can also help achieve these goals together. Accordingly, this paper considers how to package policies with one another to support local public agencies as they navigate both sides of the coin discussed previously.

Local policy actions to position agencies for thoughtful AI adoption and deployment are primarily oriented around strategic personnel and departmental coordination as a pre-condition for more systemic work. Clarifying not only which departments but also the particular staff members that will serve as a lead for AI deployment is an essential step. Importantly, this staff member should not be tasked with deciding upon where AI tools should be explored, but rather as a facilitator and resource for the necessary inter-departmental work to meaningfully incorporate AI tools into workflows.

An AI policy lead within a jurisdiction can assume responsibility for ensuring that guidelines and resources are disseminated across departments, facilitating trainings for other staff, and managing discussions with business units that are considering deployment of AI tools. An AI policy lead can also serve as the conduit for their agency and partners or knowledge from outside the jurisdiction, developing a familiarity with various potential use cases, examples of agencies that have explored or tested such applications of AI tools, and even awareness of existing tools that might be available. While an individual staff member on their own cannot ensure their agency's successful navigation of AI adoption, it is unlikely that a jurisdiction can undertake any other important policy efforts without this foundational personnel allocation.

As an agency identifies a staff member to serve as an AI policy lead, they then can commence more systemic local policy actions, such as establishing an ongoing training series for other agency staff and identifying an inter-departmental venue to discuss AI adoption. Because agency staff are already using AI tools in unsanctioned manners, agencies should offer some kind of training on the use of AI to its staff, whether or not an agency chooses to build or procure other AI tools. Agencies should consider implementing a dual-track training regimen of the kind used in the City of San Jose: one required training for all staff and another more in-depth training series for those seeking to become AI leads within their own departments.

These trainings can solve for two distinct needs at local agencies. First, all staff should be generally aware of how AI tools work, what they cannot do, and be advised on potential concerns about data safety and privacy. This more general training can help staff consider where in their workflows the use of AI might prove helpful. These general trainings should also be utilized as venues where agency leadership address staff anxieties about automation or displacement. By more directly engaging all staff around AI, rather than avoiding the issue or developing deployment plans within smaller groups, leadership can foster greater trust that is important for any adoption or eventual utilization of AI tools.

The other training series envisioned is intended for staff who are either nominated by their unit's leadership or volunteer to participate in more in-depth training(s) for staff members who become catalysts for workflow innovation. This training track is intended to help departments cultivate staff who are motivated to support exploration of new ways of incorporating AI tools and might later present ideas for how to optimize or reorganize workflows, as well as provide subject matter expert level feedback if a tool is identified for consideration within their business unit. These early adopters also become resources to their colleagues as issues or opportunities present themselves in use of official or unsanctioned AI tools. A local agency's approach to AI training should both set a "floor" on general staff understanding of AI functions as well as raise the "ceiling" of analytical capabilities of staff identified across business units who can help serve as catalysts for innovation.

Concurrently with establishing a training protocol, jurisdictions should empower their lead AI policy staff to issue guidelines on how AI can or should be utilized as well as identify an inter-departmental body to navigate related questions. Guidelines on AI use by the jurisdiction serve several important functions. An official statement on the agency's approach to AI has several audiences it can address, but particularly residents and staff. Such guidelines become a reflection of the jurisdiction's guiding principles as it relates to how AI tools might be integrated in staff workflows or contact points with residents, with an intention of conveying forethought and transparency in how important questions in deployment are navigated. The development of such guidelines itself can become a process for addressing potential concerns among residents or staff. Agencies might consider presenting a draft version of an AI adoption framework

at a public meeting to solicit feedback from elected leaders and the public before implementing this as practical guidelines to be shared among staff.

California can empower this work within local agencies by providing resources from a state level. Though peer-to-peer resources are increasingly available through the GovAI Coalition and several organizations provide technical assistance or general recommendations for agencies exploring adoption of AI tools, there are benefits from official engagement by the state that cannot be replicated in those settings. A state agency can reach local jurisdictions at scale, collect and report data or findings from across the state, and signal credibility (or at least permissibility) to an approach or tool.

Interview participants from both local and state agencies envisioned a constructive role for trainings facilitated by a state agency, as well as a community of practice to bring together local jurisdictions. Both these means for developing local staff capacity could be organized either by function (i.e., for general staff, AI policy leads, procurement managers, CIOs, etc.) or by use case (i.e., benefits processing/delivery, public works, public safety, etc.). Another valuable role that a state agency can play is collecting data and feedback from local agencies across California. One such opportunity presented in an interview was for the state to survey local agencies around how staff time is utilized to help identify both where AI tools might be the most impactful (in addressing tasks that command a large percentage of staff time) and where some jurisdictions might have found tools that are helpful (by identifying agencies with significantly lower shares of staff time applied toward time-intensive tasks). Through its reach and the state's inherent role as an authority relative to local jurisdictions, a California agency can play a key role in building capacity and sharing information.

A California state agency can play another important role in gathering and sharing information with local agencies through its interactions with or regulation of vendors of AI tools. Given that voluntary model cards and tool information sheets are used in a range of manners by vendors, California should consider how its state agency might play a role in helping promote greater standardization of disclosures. Agency leaders or policymakers can consider two different approaches to standardizing information made available by vendors,

either through administrative action to create a voluntary, “opt-in” for vendors or legislation to empower agencies to enforce requirements on disclosures. Administrative action could take one form by making tools available to local agencies for testing purposes through a “sandbox” that allows testing with synthetic datasets so potential users can explore their applicability.

A statewide approach to AI adoption among local agencies should also consider data sovereignty and storage costs. California should begin treating residents’ data like a public good, rather than a cost driver. This can enable more confident utilization of data for efforts to improve service delivery, while also better protecting the public’s data. This shift requires working with CDT to not only establish new rate structures for data storage services for local and state agencies but also the kind of infrastructure required for AI tools specifically. As part of this effort, California can target appropriations for CDT to offer a discounted rate for data that contains PII and state officials reasonably believe could be subject to potential intrusion.

Rather than making data with high rates of egress and sensitive PII the most costly for public agencies to store, California should subsidize its storage so it can be more safely stored by a state entity and remain available for operational uses. If data were held by CDT, it could also become an important asset for training or fine-tuning AI tools intended for public sector uses that could be made available as open-source offerings to California public agencies. A California compute center not only can address a significant cost for local agencies, but can help ensure data sovereignty and facilitate the proliferation of AI tools built by and for public agencies to enhance their service delivery.

What policy actions today can guide AI deployment toward better impacts tomorrow?

The daily advancements of AI capabilities can complicate long-term planning around how deployment within organizations should proceed. As California local and state policymakers consider how to refine the above concepts for implementation in the near-term, they must do so with humility about the approach and the need for its iteration. The well-recognized lack of existing staff capacity in this space also suggests certain near-term activities while cautioning against other, more rigid approaches. Despite these caveats, there are still several immediate policy and administrative actions that California can pursue to align a statewide framework and resources with local innovation and governance.

As documented in Section 2, many local governments are already implementing AI governance policies, particularly as it relates to internal coordination and transparency. The policy levers for local government presented above are largely the same as the implementation steps they should pursue now. Smaller and medium size jurisdictions in particular will benefit from existing templates already available from those publicly approved or offered by the GovAI Coalition. Identifying AI policy leads or convening inter-departmental bodies to discuss data and AI governance are steps that do not require approval from a City Council or Board of Supervisors, but rather can be proactively implemented by agency leadership. Staffing and guidelines alone are insufficient without forging broader buy-in among leadership and frontline staff. Creating buy-in is a critical means of maintaining the agency's focus and also engendering trust for the change management process that comes with the adoption of new tools.

Local governments can also implement strategies to engage their residents in this process and reinforce the public trust. Conducting an audit of existing AI systems in agency operations and publishing an inventory of such uses can serve two goals. First, this can help catalyze important internal conversations about these tools and the processes in which they are embedded and build experience for this kind of future work from examples with which staff are familiar because they are already in deployment. Second, this process can also help convey to the

public that the agency will be transparent in how it explores any new adoption of AI tools moving forward. Again, public agencies will benefit from a number of examples of inventories that are already available in other jurisdictions, as well as the GovAI Coalition. As a local agency does move forward with other new AI tools, it might also consider creating a means for resident engagement in the process, whether a "nomination" type process of the kind being piloted in Long Beach or through public events like hackathons where residents can work directly with agency staff as they explore means of leveraging data and new tools to address public problems. All of these strategies develop greater internal awareness and capacity while also engaging residents in the process.

While local governments have already demonstrated how to start implementing policies around AI adoption and deployment, California has yet to take any comprehensive state actions in support. In the immediate term, California could begin by extending existing spaces to local agencies, potentially in partnership with entities like the GovAI Coalition. Existing CDT training for staff or ODI's community of practice can be expanded upon and tailored to suit the needs of local agencies.

Another valuable and low-cost activity for the state to lead would be in surveying local agencies about their uses of staff time, which can help identify which functions would be best suited to address at a statewide level through trainings or collaborative work to identify potential AI tools in support. Though there might be some cost associated with its staffing, a California agency like CDT or ODI should develop a sandbox for vendors that would like to voluntarily make their tools available for testing with a synthetic dataset. Local agencies are beginning to utilize this approach already and can be engaged as partners to help scale this practice and make it available to less-resourced local agencies. These initial actions are relatively lower cost, but will require staff time. Wherever possible, CDT, ODI, or any other involved state agency should seek to partner with the growing community of local practitioners and technical assistance entities; where California agencies can provide the most value in these kinds of activities is in their reach and the imprimatur of the state's involvement.

State legislators can also begin to take more focused action through introducing proposals to require that CDT develop standardized tool disclosure templates

that local and state agencies must utilize when procuring AI tools. Rather than the legislature attempting to incorporate all the necessary informational points required of such disclosures within a bill, this level of detail should be entrusted to CDT to navigate in partnership with organizations like the GovAI Coalition and in conversation with industry representatives. By investing this authority with a state agency, rather than embedding too much detail within the legislation itself, the policy can be adaptive to update disclosure templates if new model functionalities require different kinds of disclosures. However, legislation can and should articulate clearly the kinds of elements that a standardized disclosure template should address. This kind of state action can strike a balance in providing agencies with the additional information on vendor-provided tools they seek, while also standardizing the information that vendors provide and reducing the friction they encounter when working across agencies. If framed appropriately and stakeholders from both the public and private sector are engaged, as well as partners at CDT, an entrepreneurial legislator may be able to forge a strong coalition to see through such a bill.

Construction of a state compute center is not entirely without precedent given CDT's existing service offerings and implementation can be informed from this experience. Additionally, scholars have analyzed the potential fiscal implications of constructing national level compute centers and how that might defray costs for federal agencies.

Los Angeles is the most heavily populated city in California with more than 3.8 million residents, twice the size of San Diego, the next largest (United States Census Bureau, 2025). The City of Los Angeles executed a three year contract with CDT for approximately $10.5 million in 2019 to cover data storage costs and related staff services (Eidam, 2019). While data collection, use, and storage practices vary significantly across jurisdictions, this contract might help indicate an approximate "per resident" cost to data storage at CDT to help contextualize what amount of funding might be required to subsidize use of the state storage option. In 2025 dollars, Los Angeles expended $1.16 per resident annually in contracting for CDT's data storage and related staff services. While the costs for conversion of a municipal mainframe to CDT storage is not directly comparable with costs associated with a state compute center suitable for running AI tools at scale, it does help establish a floor for projections.

When the Little Hoover Commission explored this topic, leading national expert on public AI infrastructure testified that construction of a public compute center would yield significant savings relative to purchasing cloud credits with private sector vendors; at a national level, public infrastructure would be seven times cheaper than purchasing credits, though a state-level compute center or consortium would likely not yield the same degree of savings (Wald, 2024). Though smaller or more fiscally impacted agencies might be less resourced to purchase compute from CDT, providing subsidies for storage associated with PII can help further defray costs. Because local agencies are already bearing costs for these services, and there is a record of shifting to CDT services, California should not let financing concerns alone stall this effort.

While financing should not be a barrier to taking on the project of a state compute center, it must still be approached cautiously given the state's immediate fiscal challenges. To feasibly take on this project, CDT should start with a pilot effort to construct the kinds of data infrastructure necessary for the use of AI tools. The capital outlays required to purchase the necessary equipment for this service lend themselves to a one-time appropriation that can also be considered in the context of the fees that state or local agencies might pay to access such services. Because CDT can anticipate that at least some portion of the costs of providing this service will reasonably be provided by agency customers, even if paying at a lower rate than private sector providers, CDT can finance capital costs against future revenue. CDT should proceed conservatively, starting with a state investment to develop the pilot at a modest level while studying the potential interest in scaling and fee structure focused on cost recovery. As part of this study, CDT should also present policymakers with options for what a subsidy program for PII data might entail in terms of annual appropriations.

As a compute center develops and matures, it can also serve as a resource for datasets for public sector developers. The scale of data stored at the compute center would make it a valuable development asset itself. Hosting a sandbox for vendor tools as outlined in this section would be even more easily facilitated with a compute center as it would have a rich starting point from which to derive synthetic datasets across a variety of domains. It is also likely that in the years required to scale a state compute center, the capabilities of platforms that

support coding and creation of digital or AI tools will continue to accelerate and facilitate an even greater rate of application development within state and local agencies. There are already instances where local agencies have tested and deployed open-source models to advance their work. These models can be housed at a compute center and made available to public agencies, but potentially even trained on a much larger, more diverse, and better protected dataset at the compute center. In time, a compute center could serve as a shared resource benefiting developers from both the public and private sectors.

Summary of Recommendations

Presented below is a summary of the most urgent policy and administrative actions that policymakers and public agency leaders should implement to position California to strategically manage its adoption of AI tools for constructive innovation in service delivery:

Local policy & administrative actions

- Designate staff to serve as "AI Policy leads" and help coordinate interdepartmental efforts
- Create an internal training program to address the needs of both general staff that might interact with general purpose AI tools and "change catalysts" from across departments that are interested in helping identify specific instances where AI tools might be helpful
- Conduct an audit of current uses of AI tools and datasets upon which they rely or might be utilized in the future
- Draft and publish internal guidelines on how the agency will utilize AI tools
- Publish a public-facing inventory of AI uses
- Review and update procurement guidelines to ensure they reflect questions and issues specific to AI tools when appropriate
- Engage peer jurisdictions to understand where they have found constructive uses of AI and facilitate interagency knowledge transfer on practices like piggy-back contracting or sandbox testing

State policy & administrative actions

- Require that CDT issue guidelines on information that public agencies must receive when procuring AI tools from an external vendor
- Provide funding to CDT to expand training efforts and initiate communities of practice that are oriented toward local agencies
- Provide funding to CDT to facilitate a "sandbox" program which local agencies can test AI tools from vendors that volunteer to make their products available to use with synthetic datasets
- Initiate a pilot program in which CDT offers data storage and compute necessary for AI tools available to local public agencies

- Update data protection laws for data held by the state on behalf of local agencies
- Conduct a cost feasibility study on how CDT could scale its data storage and compute pilot program, including potential financing via revenue bond

Charting a path toward greater public impact in local government

In looking at how to chart a path toward leveraging technology tools for better, more resilient public services from local government, California faces both a significant opportunity and an urgent moment. AI tools are already being used across local governments, but without sufficient policies or resources that can help prompt deliberation about how they are developed, tested, or overseen. Conversely, local governments are looking for means to better serve residents now in a climate of significant fiscal headwinds. California's policymakers, both locally and at a state-level, should act quickly to respond to these parallel opportunities and challenges now while also anticipating the continued extension of the capabilities and breadth of AI tools available to public agencies.

The above recommendations are presented as a coordinated approach for how state policymakers and local agency leaders can walk this path together. Failure to account for action at either level will impede how the state progresses in its ability to strategically deploy technology tools that benefit shared constituents. State disclosure frameworks and resources are only valuable to the extent they are used by local agency staff who are well-trained in understanding the core technologies and have thoughtfully engaged around the specifics of a use case. Local leaders can invest significant time in creating their own governance policies and staff coordination mechanisms, but may still confront challenges seeking consistent information from vendors or safe, cost-effective data storage and compute. These recommendations are also intended to scaffold support by providing resources and greater access to information now while building toward new forms of support later through potential expanded roles for the compute center.

At this stage, California is navigating a course toward innovation in government operations utilizing these new technologies. As public agencies and policymakers move forward together, an inevitable next leg of the journey will be exploring what such advancements mean for the underlying logic of state-funded and locally administered programs – whether innovation in those policies is necessary because of the new tools available to local agencies in delivering them.

Works Cited

Albergotti, R. (2026, February 27). “Trump Administration Bans Anthropic AI in National Security Crackdown.” Semafor. Retrieved from: https://www.semafor.com/article/02/27/2026/hours-after-pentagon-bans-anthropic-openai-strikes-defense-deal

Blackwell, B., Caprara, C., Rountree, J. Casey, P., Vanderford, D. & Battis, C. (2024, November 18). “The Homelessness Prevention Unit: A Proactive Approach to Preventing Homelessness in Los Angeles County.” California Policy Lab. Retrieved from: https://capolicylab.org/the-homelessness-prevention-unit-a-proactive-approach-to-preventing-homelessness-in-los-angeles-county/

Blackwell, B., Caprara, C., Rountree, J. Casey, P., Vanderford, D. & Battis, C. (2025, July 10). “Early Outcomes from the Los Angeles County Homelessness Prevention Unit.” California Policy Lab. Retrieved from: https://capolicylab.org/early-outcomes-from-the-los-angeles-county-homelessness-prevention-unit/

Burger Ayogu, K., Hoffnagle, E., Liebman, J., & Mertz, K. (2025) “What is Procurement Excellence?” Partners for Public Good. Retrieved from: https://partnersforpublicgood.org/procurement-excellence-network/wp-content/uploads/sites/2/2025/03/What-Is-Procurement-Excellence.pdf

California Department of Technology. (2021). “Bridge to the FutureL Statewide IT Annual Report - 2021.” Retrieved from: https://www.cdt.ca.gov/wp-content/uploads/2022/01/Annual-Report2021-Final.pdf

California Department of Technology. (2025). “Statewide Information Technology Annual Report, 2024-25.” Retrieved from: https://www.cdt.ca.gov/wp-content/uploads/2026/03/2024-2025-Annual-Report.pdf

California State Assembly Committee on Local Government. (2015, July 15). “Senate Bill 272 (Hertzberg) Committee Analysis.” Retrieved from: https://leginfo.legislature.ca.gov/

California State Assembly Committee on Public Safety (2025, July 1). “Senate Bill 524 (Arreguin) Committee Analysis.” Retrieved from: https://leginfo.legislature.ca.gov/

California Office of Data & Innovation. (2024). “Data Strategy Report 2024.” Retrieved from: https://innovation.ca.gov/reports/data-strategy-report-2024/#strategic-plan

California Office of Data & Innovation. (2026). “California Statewide Data Strategy 2026-2027.” Retrieved from: https://innovation.ca.gov/reports/california-statewide-data-strategy-2026-2027/

Chen, M., Esposito, M., Giannella, E., Guo, Z. Gosciak, J. & Koenecke, A. (2026, March 13). “Helping the Helpers: Evaluating a GenAI-powered assistive chatbot for caseworkers.” Nava Labs, Better Government Lab at Georgetown University, & Department of Information at Cornell University. Retrieved from: https://assets.ctfassets.net/t2ekr6eg3fr3/6Mw21Ui0jPkfXGW22BwUhb/9dc81c9ff2a9c251cfc9d484d071c17e/Helping_the_Helpers__Nava_Labs_March_2026.pdf

CompTIA Public Technology Institute. (2023). “2023 State of City and County IT National Survey.” Retrieved from: https://comptiacdn.azureedge.net/webcontent/docs/default-source/research-reports/10385-pti-2023-state-of-city-and-county-it-national-survey-online.pdf?sfvrsn=5b258561_0

Descant, S. (2025, November 13). “Chatbot Not Found: Long Beach Shares Lessons Learned From 'Elby’” Industry Insider (GovTech). Retrieved from: https://insider.govtech.com/california/news/chatbot-not-found-long-beach-shares-lessons-learned-from-elby

Eidam, E. (2019, March 21). “CDT Announces $10.5M Data Center Contract with Los Angeles.” GovernmentTechnology. Retrieved from: https://www.govtech.com/computing/CDT-Announces-105M-Data-Center-Contract-with-Los-Angeles.html

Feng, A., & Wu, S. (2019, May 1). “The Myth of the Impartial Machine.” Parametric Press, Issue 1. Retrieved from: https://parametric.press/issue-01/the-myth-of-the-impartial-machine/

Geraghty, L. (2021). “Cities in the Cloud.” National League of Cities Center for City Solutions. Retrieved from: https://www.nlc.org/wp-content/uploads/2022/02/CS-Cities-in-the-Cloud-Report.pdf

Goldsmith, S. & Yang, J. (2024). “AI and the Transformation of Accountability and Discretion in Urban Governance.” Retrieved from: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4968086

Goodman, J. (2025, August 5). “Big Cities Face Deficits: Should States Worry?” Pew Charitable Trusts. Retrieved from: https://www.pew.org/en/research-and-analysis/articles/2025/08/05/big-cities-face-deficits-should-states-worry

Government AI Coalition. (n.d.). “Addendum [X]: Requirements for AI Systems.” Retrieved from: https://www.sanjoseca.gov/home/showpublisheddocument/118649/638774761928370000

Government AI Coalition. (n.d.). “AI FactSheet.” Retrieved from: https://www.sanjoseca.gov/home/showpublisheddocument/118633/638774754202870000

Government AI Coalition. (n.d.), “AI Governance Handbook.” Retrieved from: https://www.sanjoseca.gov/home/showpublisheddocument/118629/638774749167870000

Government AI Coalition & Center for Scientific Evidence in Public Issues. (n.d.), “Key Questions for Government Leaders to Ask When Considering the Use of AI Systems.” https://www.aaas.org/sites/default/files/2024-11/EPICenterAIQs.pdf?utm_medium=display&utm_source=epi_center-website&utm_campaign=epi2024_aiqs&utm_id=recAwnv42v13cTFZY

Heifetz, R. & Linsky, M. (2017). *Leadership on the Line: Staying Alive Through the Dangers of Leading (Updated ed.).* Harvard Business Review.

Industry Insider (GovTech). (2025, November 26). Retrieved from: https://insider.govtech.com/california/news/californias-most-digital-cities-revealed-part-iii#:~:text=The%20city%20has%20implemented%20an,speed%20access%20for%20city%20staff.

Johnson, K. (2025, August 25). " Three bills would protect California workers from AI management, but will costs stand in the way?" CalMatters. Retrieved from: https://calmatters.org/economy/technology/2025/08/california-ai-employment-legislation/

Jones, J. (2025, November 10). "U.S. Trust in Government Depends Upon Party Control." Gallup. Retrieved from: https://news.gallup.com/poll/697421/trust-government-depends-upon-party-control.aspx

Kleinberg, J., Ludwig, J., Mullainathan, S. & Obermeyer, Z. (2015). "Prediction Policy Problems." American Economic Review: Papers & Proceedings, Vol. 105, No. 5, pp. 491–495. Retrieved from: https://www.cs.cornell.edu/home/kleinber/aer15-prediction.pdf

Lau, T. (2020, April 1). "Predictive Policing Explained." Brennan Center for Justice. Retrieved from: https://www.brennancenter.org/our-work/research-reports/predictive-policing-explained

Legislative Analyst's Office. (2022, February 22). "The 2022-23 Budget California Department of Technology." Retrieved from: https://lao.ca.gov/Publications/Report/4552

Little Hoover Commission. (2024, December). "Artificial Intelligence and California State Government." Retrieved from: https://lhc.ca.gov/report/artificial-intelligence-and-state-government/

LYT.speed. (n.d.). "Case Study: The VTA Pilot Deployment." Retrieved from: https://hs-5754551.f.hubspotemail.net/hubfs/5754551/Case%20Studies/LYT.speed%20VTA%20Case%20Study.pdf?utm_medium=email&_hsmi=353629102&utm_content=353629102&utm_source=hs_automation

Mulligan, D. & Bamberger, K. (2019, October 15). " Procurement As Policy: Administrative Process for Machine Learning." Berkeley Technology Law Journal, Vol. 34. Retrieved from: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3464203

Nabatchi, T. & Blomgren Amsler, L. (2014). "Direct Public Engagement in Local Government." The American Review of Public Administration. Retrieved from: https://digitalcommons.hamline.edu/cgi/viewcontent.cgi?referer=&httpsredir=1&article=1009&context=dri_symposia

National Association of State Chief Information Officers. (2025). "The 2025 State CIO Survey: Leading Change Through Uncertain Times." Retrieved from: https://www.nascio.org/resource-center/the-2025-state-cio-survey/

National League of Cities. (2016). “Cities 101: Delegation of Power.” Retrieved from: https://www.nlc.org/resource/cities-101-delegation-of-power/

National League of Cities. (2024, November 13). “AI in Cities - Report & Toolkit.” Retrieved from: https://www.nlc.org/resource/ai-report-and-toolkit/?

New York State Assembly. (2025). “Assembly Bill A8295D.” Retrieved from: https://nyassembly.gov/leg/

New York Information Technology Services. (2025, March 4). “New York State Information Technology Guideline No: NYS-G25-001 - IT Guideline: AI System Inventory Guidance Document.” Retrieved from: https://its.ny.gov/system/files/documents/2025/03/nys-g25-001-ai-inventory-guidance-document.pdf

OECD. (2025). “Governing with Artificial Intelligence: The State of Play and Way Forward in Core Government Functions.” Retrieved from: https://doi.org/10.1787/795de142-en.

Office of Governor Gavin Newsom. (2023, September 6). “Governor Newsom signs executive order to prepare California for the progress of artificial intelligence.” Retrieved from: https://www.gov.ca.gov/2023/09/06/governor-newsom-signs-executive-order-to-prepare-california-for-the-progress-of-artificial-intelligence/

Office of Governor Gavin Newsom. (2025, April 30). “Governor Newsom announces launch of new AI tool to supercharge the approval of building permits and speed recovery from Los Angeles Fires.” Retrieved from: https://www.gov.ca.gov/2025/04/30/governor-newsom-announces-launch-of-new-ai-tool-to-supercharge-the-approval-of-building-permits-and-speed-recovery-from-los-angeles-fires/

Office of Governor Gavin Newsom. (2025, December 16). “Governor Newsom launches new initiatives to partner with tech policy experts and accelerate responsible AI in state government.” Retrieved from: https://www.gov.ca.gov/2025/12/16/governor-newsom-launches-new-initiatives-to-partner-with-tech-policy-experts-and-accelerate-responsible-ai-in-state-government/

Office of Governor Gavin Newsom. (2026, March 30). “ As Trump rolls back protections, Governor Newsom signs first-of-its-kind executive order to strengthen AI protections and responsible use.” Retrieved from: https://www.gov.ca.gov/2026/03/30/as-trump-rolls-back-protections-governor-newsom-signs-first-of-its-kind-executive-order-to-strengthen-ai-protections-and-responsible-use

Rittel, H. & Webber, M. (1973). “Dilemmas in a General Theory of Planning.” Policy Sciences, Vol. 4, No. 2, p. 155-169. Retrieved from: https://www.jstor.org/stable/4531523

Texas Department of Information Resources. (2026, March 23). Public Sector AI Systems Advisory Board Meeting. Retrieved from: https://www.youtube.com/watch?v=uulJtspUrW0

Texas Legislature Online. (2025). “Senate Bill 1964.” Retrieved from: https://capitol.texas.gov/

Texas Legislature Online. (2025). “House Bill 2818” Retrieved from: https://capitol.texas.gov/

United States Bureau of Labor Statistics. (2025, January 28). “Union Members Summary.” Retrieved from: https://www.bls.gov/news.release/union2.nr0.htm#:~:text=The%20union%20membership%20rate%20continued,educational%20services%20(13.2%20percent)

United States Census Bureau. (2025). “Population Estimates” Retrieved from: https://www.census.gov/quickfacts/fact/table/CA/

Wald, R. (2024, May 23). Testimony before the Little Hoover Commission: Hearing on Artificial Intelligence and California State Government [Video]. California's Little Hoover Commission. Retrieved from: https://www.youtube.com/watch?v=ltFnwjWywPo

Xiang, A., & Raji, I. D. (2019). “On the Legal Compatibility of Fairness Definitions” (Version 1). arXiv. Retrieved from: https://doi.org/10.48550/ARXIV.1912.00761