



THE CITY OF ATLANTA ARTIFICIAL INTELLIGENCE COMMISSION

Final Report
May 2026



TABLE OF CONTENTS

Note from the Chair	Page 3
Purpose of the Commission	Page 4
Commission Overview	Page 5
Commission Meetings	Page 7
Recommendations	Page 10
The Path Forward	Page 15
Appendix	Page 16

NOTE FROM THE CHAIR

Last Spring, the Atlanta City Council established the Artificial Intelligence Commission, bringing together an incredible group of Atlantans – renowned experts in their respective fields – to explore and examine the growing role of artificial intelligence across government operations, public services, economic development, and community impact.

Over the course of 12 months, the Commission met publicly, reviewed emerging industry standards and policy approaches from across the country, and gathered input from technology leaders, academic institutions, community stakeholders, and City departments.

The completion of the Commission’s work marks an important step forward for Atlanta as we prepare for the opportunities and challenges presented by artificial intelligence.

These recommendations reflect months of thoughtful collaboration and public engagement focused on ensuring innovation is balanced with accountability, transparency, and the public interest. They focus on establishing governance standards for AI usage, promoting ethical safeguards, supporting workforce preparedness, and encouraging innovation that benefits residents and businesses.

They provide guidance for how Atlanta can lead responsibly in the rapidly evolving landscape of artificial intelligence. The work completed by the Commission positions the City to approach emerging technologies thoughtfully while maintaining public trust and ensuring these tools serve our communities effectively.

I want to express my gratitude for the time, thoughtfulness and passion my fellow commissioners have dedicated to this process—and to my former Council colleague Amir Farokhi for launching this initiative in the first place.

I look forward to presenting these recommendations to my Council colleagues, and working together to review them for potential legislative and administrative action—in coordination with City leadership, stakeholders, and relevant departments.

Sincerely,

A handwritten signature in blue ink, appearing to read 'M Westmoreland', with a long horizontal flourish extending to the right.

Councilmember Matt Westmoreland

PURPOSE OF THE COMMISSION

The purpose of the Artificial Intelligence Commission was to explore how to best and most efficiently deploy artificial intelligence to improve city efficiency, analyze data, and bolster constituent services.

COMMISSION OVERVIEW

In December 2024, the Atlanta City Council adopted 24-R-4521, establishing the Artificial Intelligence Commission for the purpose of “exploring the benefits, opportunities, and challenges of the use of artificial intelligence in city government.”

The Commission brought together thought leaders and subject-matter experts on a monthly basis to identify ways in which AI can enhance and improve municipal government.

Serving as members of the City of Atlanta Artificial Intelligence Commission were:

Charlotte Alexander | Professor, Georgia Institute of Technology

Donald Beamer, Jr. | Senior Technology Advisor, City of Atlanta Mayor’s Office

Nikhil Despande | Chief Digital & Artificial Intelligence Officer, Georgia Technology Authority

Mathew Garver | CEO, HiQo Solutions

Dr. Joy Harris | Director, Entrepreneurship and Innovation Institute, Georgia State University

Dr. Brandeis Marshall | Founder & CEO, DataedX Group

Jason Sankey | Chief Information Officer, City of Atlanta

Tyler Scriven | Co-Founder & CEO, Saltbox

Erika Smith | Global Partnerships Lead, Microsoft

Dr. Joseph Sutherland | Director, Emory University Center for Artificial Intelligence Learning

Larry Williams | President & CEO of the Technology Association of Georgia

John Yates | Partner, Gunderson Dettmer

Chair: Matt Westmoreland | Atlanta City Council, Post 2 At-Large

Serving as lead technical advisor to the Commission was Hyun Shin, Senior Vice President at Ohio River South

This report authored by Graham Younger, Director of Policy for CM Westmoreland and Hyun Shin

Over the course of the last year, the Commission covered a wide variety of topics, including:

- **Establishing AI Governance and Policy:** The Commission dedicated significant time to discussions on ethics, responsible AI use, and governance. We reviewed legal consequences and the latest trends in AI regulation, including the federal government's December 11th executive order and state bills, and examined best practices from other municipalities, such as the City of Boston's AI Guidelines, to inform ethical governance recommendations.
- **Assessing Current and Future Operations:** The Commission received detailed briefings on the City of Atlanta's existing AI deployments, which include AI-enabled tools like CoPilot, and Leak Detection Pilot for Watershed Management. Discussions were centered on how to use AI to improve city efficiency, analyze public data, ease car traffic, and bolster constituent services.
- **Prioritizing Workforce Readiness:** Substantial time was dedicated to understanding the barriers to AI adoption. The Commission explored how to prevent a technological access gap, particularly for college students. The group discussed strategies for developing a clear strategy for the city workforce, building an AI-fluent culture, and defining "readiness" in organizational terms.
- **Developing Procurement and Best Practices:** The Commission took a deep dive into the economics of AI and procurement processes to ensure the City avoids contracts with AI-powered products that could quickly become irrelevant. We also examined innovative, data-driven projects, including predictive algorithms used in city functions.

COMMISSION MEETINGS

Meeting 1 - May 7, 2025

An inaugural meeting was held, presided over by Councilmember Amir Farokhi. Nine members were present, with the remaining seats filled by the next meeting. Members introduced themselves and had a discussion as to the most fruitful areas for AI to be used to improve city services.

Meeting 2 - September 10, 2025

The September meeting began with an introduction of Councilmember Matt Westmoreland as the Commission's new chair, as well as an introduction of the Commission's new members. Councilmember Westmoreland then laid out a roadmap for the rest of the calendar year, scheduling presentations and speakers for October, November and December.

Our speakers were from the City of Boston: Michael Lawrence Evans (Boston's Director of Emerging Technology) and his colleague Michael Huang. They discussed Boston's AI Guidelines for generative AI and highlighted how they have attempted to change their city employee's attitudes toward GenAI use at work. Their growing edges in terms of using AI for city infrastructure included monitoring water levels to be better prepared for flooding and Project Green Light: an AI-driven project that allows traffic lights to be refreshed with data more often

Meeting 3 - October 8, 2025

The October meeting focused on the City of Atlanta's existing and future projects that deploy (or will deploy) AI technologies. CIO Sankey and Director Natalie Hall, of the Office of Digital Transformation, delivered a presentation that examined City-wide AI adoption across departments.

This included current adoption of AI-enabled programs like CoPilot, a Leak Detection Pilot for Watershed Management and multiple projects for aviation, like biometric Air Exit and employee authentication. The presentation also included proposed projects for FY26 and closed with a series of suggestions of governance recommendations, many of which will be reflected in the following pages.

Meeting 4 - November 18, 2025

The Commission's November meeting featured two speakers from our commission itself Tyler Scriven, founder of Saltbox, presented on the individual's path to AI acceptance and readiness. He detailed the psychology of learning and adoption curves and how to best build a culture with meaningful incentives and sustainable engagement to get the best out of AI.

Dr. Joy Harris also led a discussion of lessons learned at the intersection of academia and AI, and how several of the college students, especially Pell Grant recipients, fail to adopt AI and as a result, fall further behind their peers.

Meeting 5 - December 10, 2025

Our December meeting was held at Microsoft. As a city partner already equipping Atlanta with multiple AI-driven technologies, mostly critically CoPilot.

Kia Floyd, General Manager of Government Relations, introduced us to the Microsoft office space. We were presented to by Gradey Caldwell, the Account Director to the City of Atlanta, who walked us through current Microsoft tools and what the future of AI-driven products could look like over the next few years.

Trina Clark James, Community Engagement Lead, also described several of Microsoft's community outreach efforts, in particular, how to think about the shift from digital fluency to AI fluency.

Meeting 6 - January 21, 2026

The Commission heard from John Yates, partner at Gunderson Dettmer, and his colleague Frida Alim. This session focused largely on legal consequences of using AI tools, and developing practical Do's and Don't's for any given organization. They helped our group recap the latest trends in AI regulation, which included the December 11, 2025 federal executive order around AI (and its potential chilling effects) as well as bills and proposals at the state level that aim to regulate AI in Georgia and other states.

Meeting 7 - February 18, 2026

Mathew Garver presented on the economics of AI and what "readiness" means for the city of Atlanta. One focus was the degree to which AI readiness is an issue because of organizational issues as opposed to technological ones. We also took a deep dive into the current state of AI capabilities and their marginal cost.

As a result, Commission members asked specific questions as to the procurement processes and how to make certain that the City does not enter into contracts / entanglements with AI-powered products that could quickly become irrelevant.

Meeting 8 - March 18, 2026

Dr. Joseph Sutherland, of Emory University AI Learning Center presented for the group on how generative AI pilots are shaping the current business landscape. He delved into how the city website can best help improve service delivery.

His presentation contained several examples of the ways that the best businesses are using AI as well as how some businesses are using it ineffectively. The commission left this meeting with a call to develop a clear strategy in how to compete as a city workforce with AI.

Meeting 9 - April 22, 2026

Charlotte Alexander presented on how AI has been used in projects developed by Georgia Tech's Law, Data and Design Lab. These examples largely focused on the criminal justice system, but helped us spur creative thoughts of how to adapt them to city functions.

Charlotte invited Aaron Bean and Christina Callas of Asemio, a consulting firm based in Tulsa, Oklahoma, to talk through how they have used public data to support critical city functions. There was particular interest in their predictive algorithms, including one that allows cities to be better prepared to predict evictions.

Meeting 10 – May 28, 2026

The Commission met for a final time, discussed the draft final report and added four new recommendations. The recommendations were approved unanimously.

RECOMMENDATIONS

The Artificial Intelligence Commission's conversations and presentations led us to support the following recommendations:

- **Use Harm and Bias Reduction as a foundational lens for any AI deployments in the City of Atlanta**

AI systems should be evaluated and governed through a harm and bias reduction framework from project inception onward. That means requiring risk-based screening that identifies potential sources of unfair treatment, discriminatory outcomes, privacy intrusions, and safety risks; documenting mitigation strategies; and building monitoring and incident-response procedures into operational plans. Decision-makers should favor approaches that minimize automated decision-making in high-stakes contexts, require human-in-the-loop safeguards where appropriate, and mandate periodic audits to detect emergent bias or degraded performance as real-world conditions change.

- **Embed equity impact assessments into any AI procurement processes**

Every procurement that includes AI components should require a formal Equity Impact Assessment (EIA) as part of vendor selection and contract award. EIAs should assess who will be affected, potential disparate impacts across neighborhoods and demographic groups, data provenance and representativeness, and how mitigation will be measured and enforced contractually. Procurement should include clear contract language requiring vendor transparency on model design, training data sources, performance metrics disaggregated by protected class, and obligations to remediate harms discovered post-deployment.

- **Ensure that community groups have a voice in evaluation of potential harms**

Incorporate structured community engagement into project design, assessment, and post-deployment review so that residents and advocacy groups can surface lived-experience concerns and local context that may not appear in technical analyses. This can include public comment periods, advisory panels representing neighborhoods most likely to be affected, participatory workshops to review use cases and prototypes, and accessible reporting channels for ongoing feedback. Community participation should be meaningful: its findings must be documented, considered in decision-making, and used to alter or halt deployments when significant harms are identified.

- **Invest in training and change management**

Deploying AI successfully requires attention to people and processes as well as technology. Develop a citywide change management program that prepares staff and stakeholders for shifting workflows, clarifies new roles and responsibilities, and establishes requirements for when to rely on AI outputs versus human judgment. Change management should include communication plans, timelines for rollout, user support resources, and evaluation checkpoints to ensure adoption does not unintentionally reduce service quality, accountability or workforce burdens.

- **Provide role-specific training for frontline staff and related personnel**

Design and deliver targeted training tailored to the responsibilities of different employee groups-- for example, call-center agents, permitting staff, code enforcement officers, data analysts, and supervisors. Trainings should cover how underlying AI models work in practical terms, limitations and failure modes, how to interpret outputs, escalation procedures for questionable results, and privacy/security requirements. Role-based curricula make sure frontline staff can confidently use AI tools while protecting residents' rights and preserving service quality.

- **Offer hands-on learning experiences where employees can experiment with AI tools**

Complement classroom learning with safe, sandboxed environments where employees can interact with representative AI tools using synthetic or de-identified data. Practical labs, pilot projects, and simulations allow staff to see real behaviors, test edge cases, and build intuition about model reliability without risking resident data. These experiential activities accelerate adoption, reveal implementation challenges early, and inform realistic operational guidelines before citywide rollouts.

- **Continue to develop a citywide AI literacy curriculum to ensure a baseline understanding of risks, benefits and appropriate uses**

Establish a citywide AI literacy program that ensures all employees have a shared baseline knowledge of AI concepts, ethical considerations, and applicable city policies. The curriculum should be accessible (multi-modal, language-inclusive), regularly updated to reflect emerging risks and technologies, and include modules for legal/compliance staff on regulatory obligations. A common literacy baseline reduces miscommunication, supports better procurement decisions, and helps residents understand how and why AI is used.

- **Appoint Artificial Intelligence leads within key City departments**

Designate AI leads in each major department to centralize expertise and provide accountable oversight for department-specific pilots and deployments. These leads should coordinate with a central AI governance body, oversee risk assessments and pilot evaluations, serve as internal points of contact for training and change

management, and ensure that deployments comply with the city's current and future AI policy. Empowered and resourced departmental leads will help standardize best practices while adapting them to each department's mission and operational constraints.

- **Explore enhancing constituent services and community engagement**

Identify practical AI use cases that can improve constituent services and deepen community engagement while preserving transparency and equity. Examples include proactively sending residents timely notifications about permit statuses, service disruptions, or scheduled maintenance; using AI to triage requests and route cases to the correct team faster; and employing predictive analytics to reduce backlogs, prioritize repairs, or optimize traffic signal timing. All pilots should include evaluation criteria for accuracy, fairness, and resident comprehension of automated communications.

- **Prioritize workforce readiness among city residents, and explore solutions to help close the technology gap**

Advance inclusive economic opportunity by investing in workforce programs that prepare Atlanta residents for AI-adjacent jobs and digital participation. Offer community training, apprenticeships, partnerships with local colleges and nonprofits, and subsidized access to necessary hardware and connectivity. Prioritize outreach to underserved neighborhoods to close the digital divide, ensuring that benefits of automation and data-driven services are broadly shared and do not exacerbate existing inequities.

- **Set clear Key Performance Indicators to evaluate AI outcomes, especially within constituent services**

Define and publish measurable KPIs to judge the success of AI initiatives and hold deployments accountable. Core metrics should include:

- Accuracy and reliability of AI outputs, measured against validated benchmarks and monitored over time. Specific metrics should be developed around the instances of hallucinations that occur over time.
- User satisfaction with AI tools, captured through resident feedback and employee surveys to assess whether tools improve experience and efficiency.
- Equity impacts, tracked via disaggregated metrics to determine whether services become more or less accessible across demographic groups and geographies. Require regular reporting on these KPIs, and incorporate results into procurement renewals, scaling decisions, and public transparency reports so the city can learn, iterate, and halt systems that do not meet performance or equity thresholds.

- **Adopt cybersecurity and privacy requirements specific to AI systems and vendors**
Develop AI-specific security and privacy standards that apply to both internally developed and vendor-provided tools. AI systems introduce unique cybersecurity and privacy risks, including prompt injection, data leakage, model manipulation, unauthorized retention of resident data, and insecure third-party integrations. Requirements could include secure model hosting, encryption standards, logging and monitoring, restrictions on sensitive-data use, adversarial testing, and clear rules governing whether city data may be used to further train commercial models. Embedding these controls early will reduce operational and reputational risk as AI adoption scales.
- **Develop formal data governance and data quality standards for AI readiness**
Adopt enterprise-wide data governance standards that define requirements for data quality, retention, lineage, metadata documentation, access controls, and interdepartmental sharing. Many AI risks stem from fragmented, incomplete, biased, or poorly governed data. Departments should be required to assess whether datasets are representative, current, and appropriate before they are used to train, fine-tune, or operationalize AI systems. A formal data governance framework would strengthen model reliability, improve interoperability across agencies, and reduce downstream equity and compliance risks
- **Create a centralized AI system inventory and public transparency registry**
Establish a citywide inventory of all AI systems in development, piloting, procurement or operational use across departments. The registry should document each system's purpose, vendor, data sources, intended users, decision-making role, risk classification, human oversight requirements, and current evaluation status. A public-facing version of the registry can improve transparency and trust by allowing residents to understand where AI is being used and how accountability is maintained. Internally, the inventory would also help reduce duplicative procurement, standardize governance practices and support cross-department learning.
- **Establish independent auditing and periodic review requirements for high-impact AI systems**
Research requiring recurring independent audits for high-impact or resident-facing AI deployments. Audits should evaluate model accuracy, disparate impacts, cybersecurity posture, explainability, compliance with city policy, and whether systems continue to perform as intended under changing conditions. Findings should be reported to the City's leadership and summarized publicly where appropriate.

FINAL RECOMMENDATION

Creation of a permanent City of Atlanta Artificial Intelligence Advisory Board

In order to achieve the above goals and maintain the momentum of the Commission after its sunset in May 2026, we propose the establishment of a permanent Artificial Intelligence Advisory Board. This body will ensure the City of Atlanta remains connected to the expert community and continues to utilize AI in a strategic and responsible manner.

To be co-chaired by Chief Information Officer Jason Sankey and Senior Technology Advisor Donald Beamer, the Advisory Board will be composed of 8-10 members who meet quarterly (with special meetings if need be) with staffing support from the Department of Information Management and the Mayor's Office.

The Advisory Board will be assembled in Q3 2026 and will hold its inaugural quarterly meeting in Q4 2026.

Core Mandates for the Advisory Board

The newly established Advisory Board should focus on the following key areas to guide Atlanta's responsible and efficient deployment of AI:

1. **AI Policy and Governance:** Advise on and refine the City's current AI policy, reiterating a core lens of harm and bias reduction across all AI operations.
2. **Accountability and Evaluation:** Advise on clear Key Performance Indicators (KPIs) to evaluate the outcomes of AI deployments. This includes ensuring rigorous due diligence and effective procurement to avoid contracts with rapidly irrelevant AI products.
3. **Workforce Readiness and Training:** Champion investment in training and change management to develop an AI-fluent city workforce. This also includes recommending that departments where AI tools are relevant appoint a dedicated AI lead.
4. **Service Delivery and Best Practices:** Identify opportunities to leverage AI to enhance the speed and accessibility of constituent services and continuously gather best practices and knowledge from other cities and private/public sector organizations.

THE PATH FORWARD

The Commission's work has been an eye-opening experience, offering a comprehensive view of this transformative field. Our experts consistently framed AI not merely as a new technology, but as a general-purpose utility, whose full potential will be captured by the cities that understand its systemic impact.

We heard significant input on the challenges and opportunities facing the City related to city operations, constituent services, workforce, and governance. AI is rapidly becoming an entrenched part of daily life, extending far beyond specific platforms or Large Language Models (LLMs). The accelerating pace at which AI is integrating across various platforms and programming demonstrates its broad and growing impact. As Atlanta strives to become a top-five tech hub in the United States, we must continue to strategically identify how AI can make our city and its citizens more effective.

As we begin this process, our philosophical recommendation is that people must be centered in every step. The rise of artificial intelligence has also come with a rise in understandable anxiety from some segments of the public—often focused on job security or resource burden. We must constantly consider the people who may be affected by these decisions. This reinforced our understanding of the economic principle of Jevons Paradox—that as a resource becomes more efficient, people end up using more, not less of it. Successful work with AI will consistently strive to create that relationship, allowing people to build and create in ways previously impossible.

The Commission's primary practical recommendation is to create an AI Advisory Board. This body will serve to maintain a connection with the group of experts who have contributed over the past year, while developing new voices to help the city advance current policies and advise us in this consistently changing technological landscape. This group can also serve as a hub to monitor AI regulation in other municipalities and states.

We want to thank Commission members, presenters, and members of the public who have shared their expertise. We look forward to further reflection as we hone the City's expertise and continue to shape Atlanta into one of the nation's elite technology hubs.

Signed,

Donnie Beamer

Jason Sankey

Matt Westmoreland

APPENDIX

Attached: Slide Decks from the meetings (note that multiple meetings were oral presentations)

Charting the Future of City Services in Boston with AI + Machine Learning



City of Boston

September 10th, 2025
Atlanta AI Commission



Michael Huang
Product Manager
Enterprise Applications



Michael Lawrence Evans
Director
Office of Emerging Technology



Our Generative AI Guidelines (circa 2023)



City of Boston

boston.gov/ai-guidelines







City of Boston

Boston's Generative AI Guidelines v1

1. Fact check and review all content generated by AI, especially if it will be used in public communication or decision making.
2. Disclose that you have used AI to generate the content.
3. Do not share sensitive or private information in the prompts.



Our use cases keep expanding.

- Improve Productivity (e.g., Procurement) 
- Predict Floods 
- Streamline Trash Routes 
- Traffic Congestion 



Improve Productivity

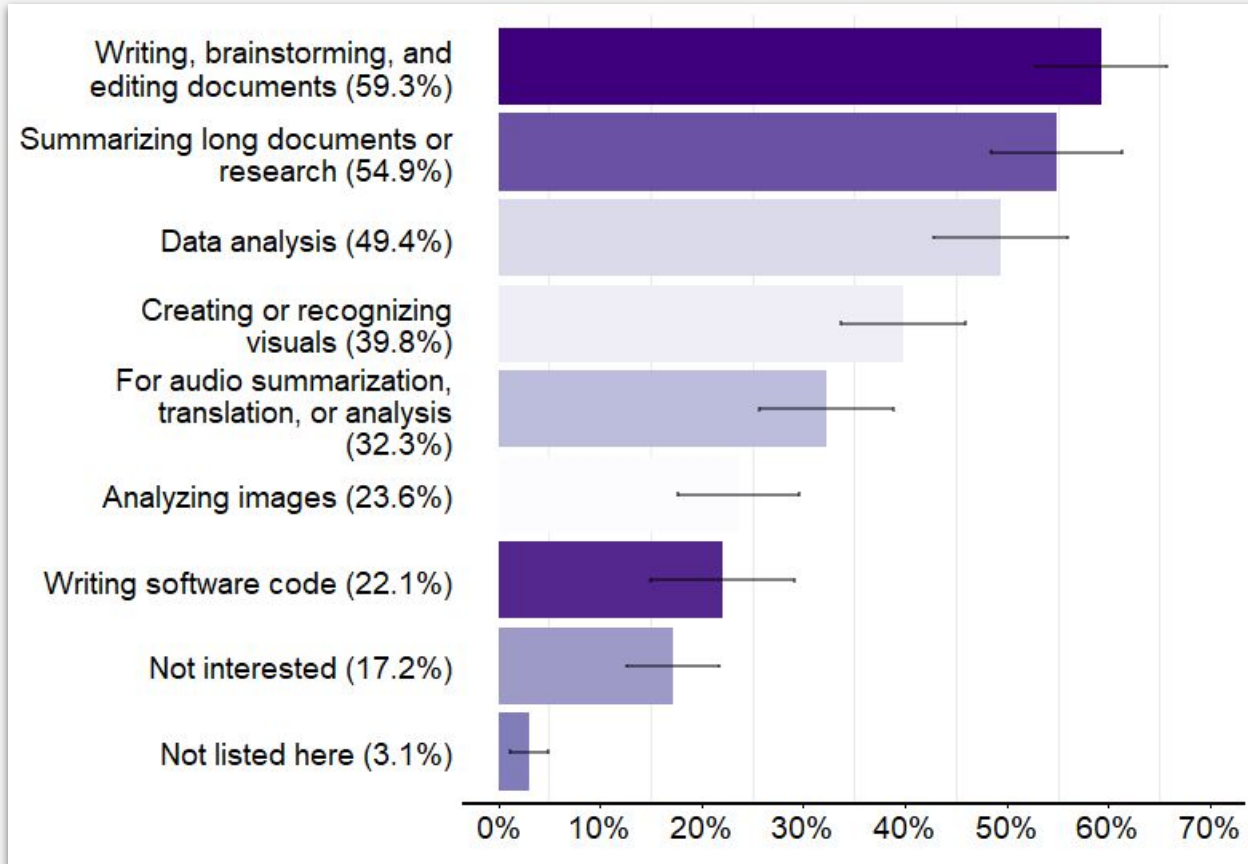


Understanding Our Community: Generative AI Survey

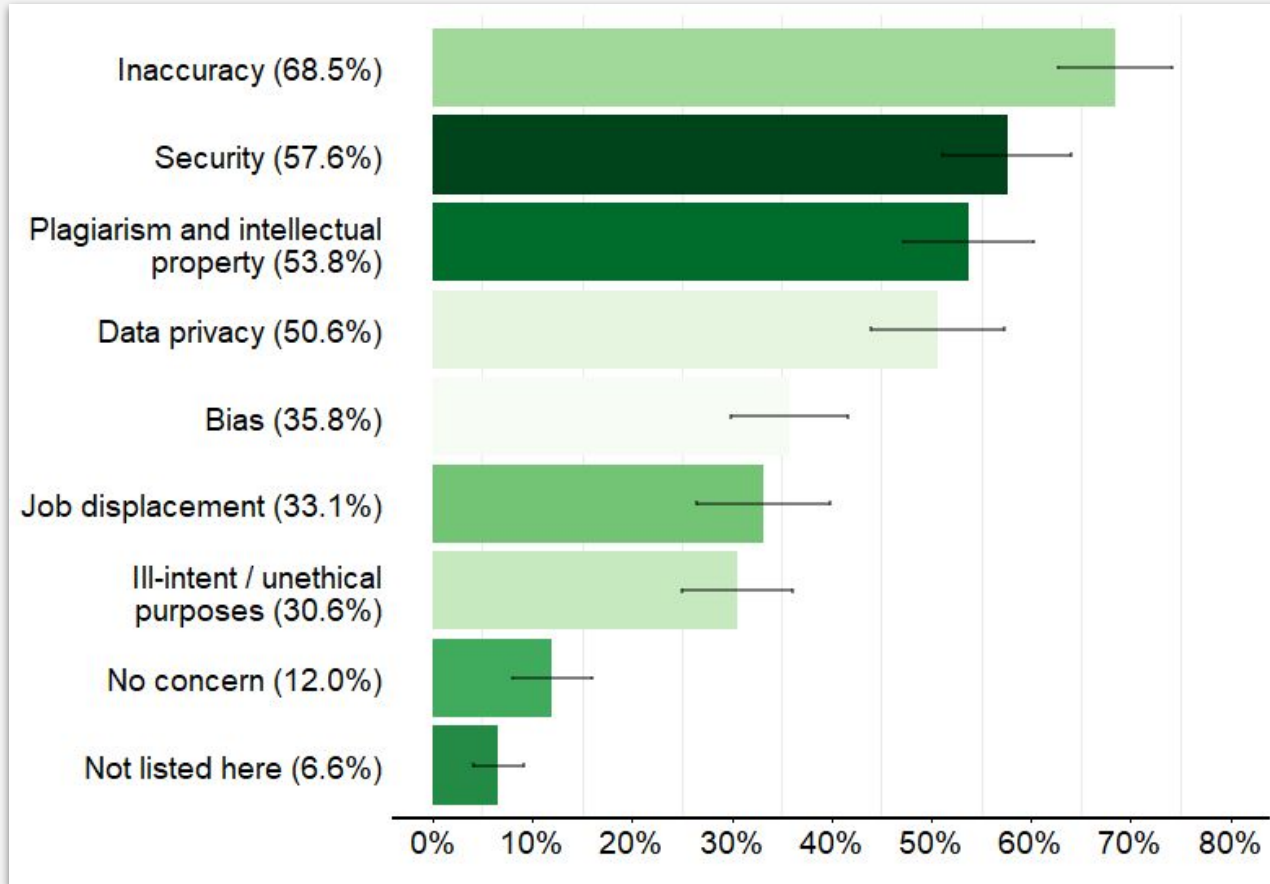


City of Boston

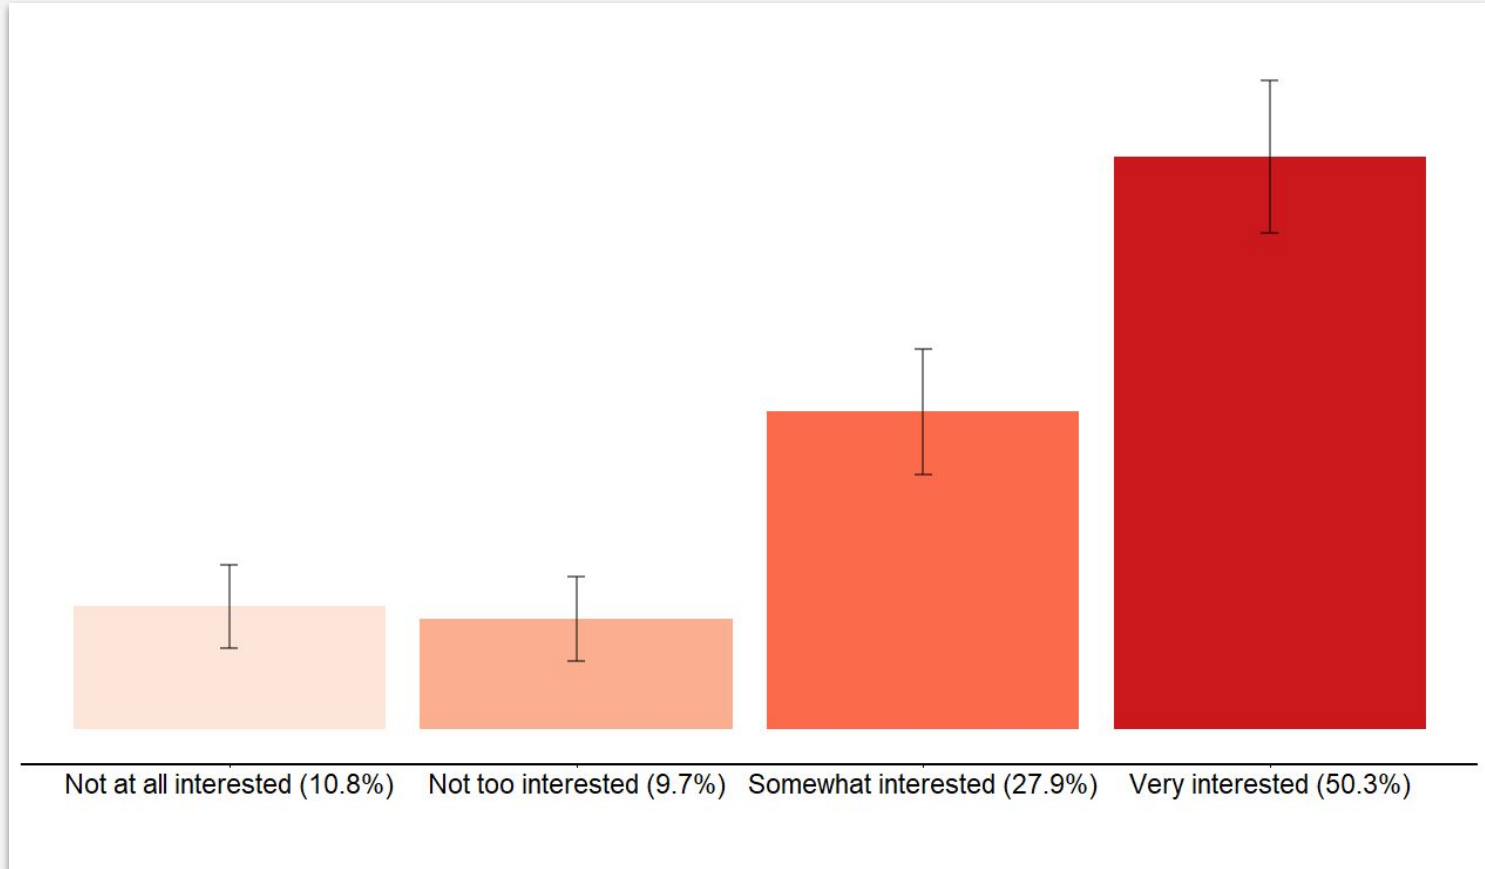
Preferred Uses of Generative AI at Work



Concerns About Using Generative AI at Work



Interest in Learning More About How to Use GenAI at Work




innovate(us)



The screenshot shows the landing page for the course 'Using Generative AI at Work'. At the top left is the 'B' logo for Boston, and at the top right is the 'innovate(us)' logo. Below the logos, the text 'ABOUT THE COURSE' is centered. The main title 'Using Generative AI at Work' is prominently displayed. Underneath the title, a short description reads: 'This free, self-paced online course helps you integrate GenAI tools into your government work, enhancing public service delivery while managing risks responsibly.' A blue 'Start Learning' button is positioned below the text. At the bottom of the page, there are two video thumbnails. The left thumbnail shows a slide titled 'Responsible AI for Public Professionals Using Generative AI at Work' with a '01:55' duration. The right thumbnail shows a man in a green sweater in an office setting, with a 'CITY of BOSTON' logo in the top right corner and a '02:33' duration.

- **GenAI basics**
 - Prompts
 - Ethical implications
 - Hallucinations/Bias
- **Go at your own pace**
- **Certification process**





Identifying Personas

Communicator

- Brainstorming
- Summarization
- Draft creation




Analyst

- Co-pilot
- Insights
- Code generation
- Process improvements




Identified Use Cases



Analytics of Usage

Daily Active Users vs. Date



Gmail - More Generation

Docs - More Summarization

Slides - More Summarization

Sheets - More Summarization



Amazon Bedrock

BROAD CHOICE OF MODELS

AI21 labs

Effective reasoning & rapid analysis for long context windows

JAMBA

amazon

Frontier multimodal intelligence at low-latency, Agent & RAG Applications, high-quality image & video generation

AMAZON NOVA

ANTHROPIC

Advanced reasoning & coding capabilities, including computer use skills

CLAUDE

cohere

Multimodal search & advanced retrieval powering multilingual knowledge agents

COMMAND
EMBED
RERANK

Luma

High-quality video generation from text & images

LUMA RAY 2

Meta

Advanced image & language reasoning

LLAMA

MISTRAL AI

Knowledge summarization, expert agents, & code completion

MISTRAL
MIXTRAL

poolside

Software engineering AI for large enterprises

MALIBU
POINT

stability.ai

High-quality AI image generation, easily deployable at scale

STABLE DIFFUSION
STABLE IMAGE

Coming soon



Launchpad

CITY of BOSTON

City of Boston GenAI Launchpad


Empowering Boston City Employees: AI-Driven Information and Support

This internal AI chatbot is designed to provide City of Boston employees with quick access to resources and support. This is an internal tool for the City of Boston. Data is collected to better understand and expand GenAI use cases.

[Get Started](#)


Featured Applications

[Email Writer](#)




Refine your emails to be clearer and more professional

[Language Translator](#)




Expertly translate and understand other languages

[Agenda Planner](#)




Create clear meeting agendas for your team

[Policy Consultant](#)




Explore, interpret, and summarize Massachusetts government policies

[Summary Writer](#)




Create concise summaries

[Document Summarizer](#)




Align content to your style brand

[Daily Planner](#)




Organize your day around your goals

[Memo Writer](#)



Write clear, concise, and professional memos

[Tone Tuner](#)



Adjust your message to your desired tone

Additional Resources

Home

► Recent Sessions

▼ Applications

- Email Writer
- Language Translator
- Agenda Planner
- Policy Consultant
- Summary Writer
- Document Formatter
- Daily Planning
- Memo Writer
- Tone Tuner

▼ Chatbot

- Playground
- Multi-chat playground
- Sessions
- Models

▼ Retrieval-Augmented Generation (RAG)

- Dashboard
- Semantic search
- Workspaces
- Embeddings
- Engines





Next Steps - Fall 2025

Rollout

Monitor

Refine



City of Boston



AI + City Infrastructure








Project Green Light 🚦

TRAFFIC IN BOSTON

- Boston is ranked **8th in the world** for traffic delays.
- Working to manage congestion in the face of growing regional traffic:
 - **Between 2019 and 2024, daily miles driven grew by 13%**
- We can find **low barrier ways to retime traffic signals with fresh data.**
 - Signals are retimed every 5 years. It's an expensive and slow process.
 - Traffic patterns change all the time and real-world conditions don't always match the model.

**Million Miles Driver per Day
in Greater Boston**



Source: MA Vehicle Census



NO LEFT TURN


MASS

ONLY STOP ON RED

Huntington Ave and Opera Place

PROJECT GREEN LIGHT: Huntington Ave. & Opera Place

Observation



Huntington Ave. (S Rdwy) & Opera Place

[Open in Google Maps](#)

High delay time detected on **Westbound** from "Massachusetts 9" to "Huntington Avenue" between the hours **15:00–19:00** on workdays.

Recommendation

Completed

15:00–19:00 on workdays

Applies to workdays only, excluding holidays and events.




Target

Adjustment

Details

Type




Phase P1

+5 sec

Increase green time by 5 seconds

Green time



Phase P2


-5 sec

Decrease green time by 5 seconds

Green time

PROJECT GREEN LIGHT: Huntington Ave. & Opera Place

Observation



Huntington Ave. (S Rdwy) & Opera Place

[Open in Google Maps](#)



High delay time detected on **Westbound** from "Massachusetts 9" to "Huntington Avenue" between the hours **15:00–19:00** on workdays.

Recommendation Completed

15:00–19:00 on workdays


Applies to workdays only, excluding holidays and events.



Target	Adjustment	Details	Type
 Phase P1	+5 sec	Increase green time by 5 seconds	Green time
 Phase P2	-5 sec	Decrease green time by 5 seconds	Green time

MEASURING IMPACT: Huntington Ave. & Opera Place

Recommendation Summary Completed



Huntington Ave. (S Rdwy) & Opera Place

[Open in Google Maps](#)

Green time shift recommendation between the hours **15:00–19:00** on workdays.

Impact Summary

13

Seconds saved
per driver

24%


Reduction in
waiting time

57%

Reduction in
split failures

MEASURING IMPACT: Huntington Ave. & Opera Place

Recommendation Summary Completed



Huntington Ave. (S Rdwy) & Opera Place

[Open in Google Maps](#)

Green time shift recommendation between the hours **15:00–19:00** on workdays.

Impact Summary

13

Seconds saved
per driver

24%

Reduction in
waiting time

57%

Reduction in
split failures

Signal Priority for School Buses (Bloomberg Mayors Challenge)



We're learning more about:

- Environmental Impact of AI
- COB employee knowledge and training gaps in GenAI
- Retrieval Augmented Generation
 - Fine-tune AI models with your documents
- Data Privacy and LLMs (including on-device AI and running our own models)

Thank You!

Artificial Intelligence Current State

JASON SANKEY, CHIEF INFORMATION OFFICER

Prepared by: Natalie Hall
IT Director, Office of Digital
Transformation

AI Projects

The City of Atlanta Department of Information Management, in collaboration with city departments, has been an early adopter and the first municipality to deploy AI technologies.

Examples of projects deployed impacting multiple departments:

- Office of Revenue and ATL311 kiosk with avatars – utilizing Azure OpenAI: July 2024
- ATL311 website and chatbot generative AI: December 2024
- Grammarly AI: June 2025
- Microsoft Copilot for employees: January 2025
- Microsoft Copilot Chat: September 2025
- ServiceNow generative AI: October 2025



ONE SAFE CITY

APD:

- Current: Flock/Fusus, Clearview
- In Flight: Microsoft CoPilot
- FY26: Axon Draft One, Bilingual services

E911:

- In Flight: Carbyne, CommsCoach
- FY26: AI Call Handling Chatbot

AFRD:

- Current: Microsoft CoPilot
- FY26: Drones, AI VR Training, Predictive Analytics, Power BI CoPilot

Solicitor's Office:

- Current: CoPilot
- FY26: AI Case Management

Corrections:

- Current: Social Media AI, CoPilot
- In Flight: Power DMS, Inventory AI, ICE Automation
- FY26: Oracle ATL Cloud, Remote Monitoring, Predictive Case Management

Municipal Court:

- Current: ChatGPT
- In Flight: CoPilot
- FY26: Chatbot, Lexis Nexis, Courtsmart, Translation, AI Presentation Tools

Public Defender:

- Current: CoPilot, Monday.com
- FY26: Case Management, Lexis Nexis, Westlaw, Axon Draft One



A CITY OF OPPORTUNITY FOR ALL

Office of Revenue:

- Current: Avatar, CoPilot, ChatGPT, Siri
- In Flight: CoPilot
- FY26: Document Recognition

Controller:

- Current: ChatGPT
- FY26: AI in Workiva, daily ops

Budget:

- Current: ChatGPT
- FY26: SOPs, Presentations, Graphics

Admin & Legislative:

- Current: ChatGPT, Presentation Bot, Grammarly
- FY26: SOPs, Legislation, Presentations

Treasury:

- In Flight: CoPilot
- FY26: Debt Portfolio Reporting, Banking Reports



A CITY BUILT FOR THE FUTURE

Watershed Management :

- Current: Leak Detection Pilot (FIDO AI): 61 leaks found, 168M gallons saved, Sewer Line AI (CCTV-based)
- FY26: Predictive Leak Detection (330 miles), Acoustic Sewer Monitoring, enQuesta AI for customer engagement
- In Flight: AMI Pilot

Transportation (ATLDOT):

- Current: CoPilot, EsriGIS, ChatGPT
- In Flight: Fusus Integration, Drones
- FY26: Autonomous Vehicles, Signal Optimization, BRT, Digital Asset Management, Crowd Management

Aviation:

- Current: Biometric Air Exit (Delta + CBP), Employee Authentication (visual recognition), TSA CAT 2 facial recognition
- In Flight: Biometric Exit Expansion (all international gates), Employee Authentication Phase II, Enhanced Passenger Processing (EPP), Autonomous Vehicle Strategy, Internal & External Chatbots
- FY26: AI Overlay for CCTV (searchable, real-time video analytics)

Public Works:

- Current: CoPilot, Routeware, Big Belly, City Detect, Geo Tab
- In Flight: AssetWorks, Maximo, Routeware Cameras
- FY26: Wireless Fueling



EFFECTIVE & ETHICAL GOVERNMENT

Contract Compliance:

- FY26: Internal Reporting, Data Visualization, Ops Standards

Inspector General:

- Current: Grammarly/OpenAI
- FY26: Monday.com Workflow

Law:

- Current: CoPilot, Monday.com, Siri, Gemini
- In Flight: CoPilot

HR:

- Current: ChatGPT, CoPilot, Apple Intelligence
- FY26: Daily Ops Support

Audit:

- Current: ChatGPT, CoPilot
- FY26: None

Procurement:

- Current: Hazel, ChatGPT, ATLcloud AI
- In Flight: CoPilot
- FY26: Oracle AI, Daily Ops

Ethics:

- Current: CoPilot, META AI
- FY26: Ethics Bot, AI for Audits, Case Mgmt, Reports

COMMON USE CASES

- Biometric security (Aviation, Public Safety)
- Predictive maintenance (Watershed, Public Works)
- AI-driven customer service (Revenue, Courts)
- Legal research and case management (Law, Public Defender)

FY26 PRIORITIES

- Hospitality concierge
- Agentic AI POC
- Expand AI in infrastructure, public safety, and citizen services
- Integrate AI into daily operations and reporting
- Strengthen governance and cross-department collaboration
- Smart parking solutions

Governance Recommendations

- **Establish AI Governance Framework**
 - Define standards, ethics, and data privacy policies
- **Centralize AI Knowledge Sharing**
 - Create a cross-department AI working group
- **Prioritize High-Impact Use Cases**
 - Focus on automation, predictive analytics, and citizen services
- **Invest in Training & Change Management**
 - Upskill staff to work alongside AI tools
- **Monitor & Evaluate AI Outcomes**
 - Use KPIs to track ROI, efficiency, and service improvements
- **Accelerate Code, Policy, and Operating Procedures**
 - Use AI to enhance our capabilities in this space to establish a governance framework



G U N D E R S O N D E T T M E R

AI USE CASES AT THE CITY LEVEL:
PRACTICAL AND LEGAL CONSIDERATIONS

January 21, 2026

We represent *what's next.*

Meet the Presenters



Frida Alim

Data Privacy Attorney
falim@gunder.com



John Yates

Partner
jyates@gunder.com

Frida is a data privacy associate in the Strategic Transactions & Licensing group.

Frida Alim counsels clients on information privacy, cybersecurity, and innovative AI strategy, guiding high-growth technology companies and venture capital firms at the forefront of digital transformation. Her practice spans high-stakes advisory and regulatory matters for leaders in frontier AI model development, biometrics, fintech, enterprise software, autonomous vehicles, robotics, social networks, telecommunications, and entertainment technologies.

Frida advises clients on the General Data Protection Regulation (GDPR), the California Consumer Privacy Act (CCPA), the Colorado AI Act, the European Union AI Act, the Children’s Online Privacy Protection Act (COPPA), the Illinois Biometric Information Privacy Act (BIPA) and the Health Insurance Portability and Accountability (HIPAA), among other privacy, cybersecurity, marketing and consumer protection laws. Her practice includes advising on privacy, cybersecurity, and AI considerations in connection with venture capital financings, mergers, and IPOs.

Frida is a Certified Information Privacy Professional (CIPP), holding the U.S. private-sector privacy certification granted by the International Association of Privacy Professionals (IAPP), the global standard in privacy certification.

Previously, as an associate at Buckley LLP in Los Angeles, Frida advised top-tier banks, innovative fintechs, and financial services firms on mission-critical privacy, regulatory, and compliance matters.

EDUCATION

University of California, Berkeley, School of Law, J.D.
University of California, Los Angeles, B.A., Political Science

FOCUS

Data Privacy
AI & Machine Learning

ADMISSIONS

California

ACHIEVEMENTS

CIPP/US Certified



FRIDA ALIM ASSOCIATE

San Francisco
falim@gunder.com
+1 415 801 4921

John is a trusted advisor to tech companies and investors.

John is a recognized leader in the field of technology law. For more than 35 years, he has dedicated his practice exclusively to this ever-evolving field, establishing himself as a trusted advisor to countless technology companies and investors.

As a prolific author and leading voice in technology law, John is a sought-after speaker, having delivered more than 250 presentations at prestigious events hosted by organizations like the American Bar Association and the Georgia Bar Association. His expertise has been recognized by the U.S. Supreme Court, and his insights have found their way into prominent technology business and law publications.

He is a co-founder of several key organizations that foster innovation and collaboration within the technology sector, including the Southeastern Software Association, the Technology Association of Georgia, the Tech Executives Roundtable, and the Tech Leaders Forum. These efforts have earned him numerous accolades, including the Technology Association of Georgia's "Leader of Influence Award" and recognition as one of Atlanta's "Most Influential Leaders."

John's dedication to the community is evidenced by his leadership roles with the United Way of Metropolitan Atlanta, Downtown Atlanta Rotary, YMCA of Metro Atlanta, Woodruff Arts Center, and CURE Childhood Cancer.

Prior to joining Gunderson Dettmer, he was a partner at Morris Manning & Martin.

EDUCATION

Duke University, J.D.

Duke University, B.A.

- *magna cum laude*

ADMISSIONS

Georgia

FOCUS

Corporate Governance & Strategy

Mergers & Acquisitions

Fund Investments in Companies

Seed, Venture and Growth Financings



JOHN C. YATES
PARTNER

Atlanta
jyates@gunder.com
+1 404 566 7814

Agenda

- I. AI use cases and real-world challenges
- II. Practical dos/don'ts of leveraging AI tools
- III. Trends in AI regulation at the state and federal levels
- IV. Questions



Key Terms

Term	Definition
ADA	Americans with Disabilities Act.
Algorithmic Bias	When an algorithm or AI system produces results that are unfairly skewed toward or against certain people or groups, leading to discriminatory or unequal outcomes.
Chatbot	A software application that simulates a conversation with a person through text or voice, usually via a website, app, or messaging platform.
Data Minimization	Principle (sometimes reflected in privacy or AI laws) restricting data collection, use and retention of personal data to what is reasonably necessary for a specific use case.
Deep Fakes	Realistic but fabricated images, videos, or audio recordings created using artificial intelligence that make a person appear to say or do something that never actually happened.

Key Terms

Term	Definition
Digital Clone or Digital Replica	AI-generated virtual version of a real person that closely imitates their recognizable traits, such as appearance, voice, expressions, and communication style.
Disparate Impact	Discrimination that occurs when a rule or practice that appears neutral ends up disproportionately harming people in a protected group, even if there was no intent to discriminate.
FOIA	Freedom of Information Act (while FOIA is federal, for purposes of this presentation, FOIA refers generally to open-records or “sunshine” laws on the state level; these laws are intended to give the public a right to inspect or obtain copies of government records).
Human in the Loop	Refers to a system or process where a human is actively involved in supervising, reviewing, or making key decisions about what an AI or automated system does, rather than letting it operate entirely on its own.



*PART I: AI USE CASES &
REAL-WORLD CHALLENGES*

Use Case #1: Public-Facing AI – Chatbots & Service Requests

OPPORTUNITY:

- Answer routine questions (permits, billing, service requests)
- Reduce phone wait times to speak with human agent
- Route complaints to appropriate departments
- Optimize city employee workloads
- Provide immediate answers

REAL-WORLD EXAMPLES:

- Austin, Houston, and several California cities use AI chatbots for 311-type service requests
- Success: reduced response times, freed staff for complex issues. However, staff time may shift to supervising bots, handling escalation.

CHALLENGES & CRITICISM:

Accuracy/hallucination: Chatbot provides wrong information about city policy → citizen files complaint or misses deadline

Equity: Non-English speakers, elderly, or digitally disadvantaged citizens **may be pushed toward chatbot when they need human help**

Data collection: Potential for vast data collection

Accountability: Who is liable if the chatbot gives bad advice? The vendor? The city?

Best Practices: Public-Facing AI – Chatbots & Service Requests

DOS

- ✓ DO: Disclose when citizens are interacting with AI, not a human
- ✓ DO: Maintain a human escalation path for non-routine issues
- ✓ DO: Ensure chatbot works in multiple languages if your city has language-access requirements

DON'TS

- ✗ DON'T: Collect unnecessary personal data (race, health, immigration status) for the task
- ✗ DON'T: Assume chatbot accuracy; audit responses regularly

Compliance measures:



City accessibility laws (ADA, state digital accessibility)



Open records/FOIA implications (what data is retained, how long, who can request it)



Vendor data-sharing agreements (does the chatbot vendor retain/analyze city interaction data?)

Use Case #2: Hiring & Employment

Example: AI-Assisted Recruitment for City Workers (HR screening, candidate ranking)

OPPORTUNITY:

- Sort through 1,000+ applications for public works, IT, or clerical roles
- Reduce bias
- Accelerate hiring
- Reduce employee time spent sifting resumes

REAL-WORLD EXAMPLES & LAWSUITS:

- Workday lawsuit (Mobley v. Workday, ongoing): AI resume-screening tools screened out applicants based on age, race, disability; tool learned patterns from biased historical hiring data
- Amazon's hiring tool infamously learned to favor male candidates

CHALLENGES & CRITICISM:

Lack of transparency: Candidates don't know why they were rejected; can't challenge the algorithm

Algorithmic bias: If training data underrepresents women or minorities in certain roles, the **AI may reinforce historical bias**

Disparate impact: Even if not intentional, **the tool may have a racially or gender-disparate impact**, violating Title VII

Disability discrimination: Personality/cognitive tests used in **AI screening may disadvantage people with mental health or disabilities**

Best Practices: AI in Hiring and Employment

DOS

- ✓ DO: Conduct bias audit of training data and algorithm before deploying AI tool
- ✓ DO: Notify candidates that AI is being used in screening
- ✓ DO: Retain human review for all AI-recommended candidates; don't auto-reject
- ✓ DO: Keep records of AI decisions for auditing and litigation defense

DON'TS

- ✗ DON'T: Use AI as the final decision-maker
- ✗ DON'T: Use personality tests or cognitive assessments without disability-impact analysis

Compliance measures:



Title VII of the Civil Rights Act (federal; applies to city as employer)



ADA (if AI tool disadvantages people with disabilities)



State/local employment discrimination laws (many states now have explicit AI in hiring laws; Atlanta city's municipal code also prohibits privacy employers from discriminating based on criminal history and gender expression)



Public records: decisions, audit results may be FOIA-able



Union contracts (if city workforce is unionized, may need to negotiate use of AI)

Use Case #3: Policing & Surveillance AI

Example: Facial Recognition, Predictive Policing, or Algorithmic Dispatch

OPPORTUNITY

- Identify suspects faster
- Prevent crime, improve community safety
- Optimize patrol routes
- Improve safety of first responders
- Digital evidence can speed up investigations

REAL-WORLD EXAMPLES & BACKLASH:

- Facial recognition: Clearview AI sued by California activists and immigrant rights groups for scraping 3 billion photos from internet **without consent** and selling to law enforcement
- City bans: San Francisco, Oakland, Boston, Minneapolis, Portland **banned government use of facial recognition**
- Concerns: **Disproportionate impact on Black and immigrant communities**; errors in misidentification; chilling effect on free speech and political organizing

CHALLENGES & CRITICISM:

Privacy violations: Scraped data without consent; database built on stolen images

Accuracy disparities: Facial recognition is less accurate for Black women and other communities of color

Chilling effect: Activists, protesters, and marginalized communities fear surveillance

Lack of warrant/oversight: Police may use tool without judicial review

Opacity: How were the algorithms trained? What data? Who has access?

Best Practices: AI in Policing and Surveillance

DOS

- ✓ DO: Establish clear policies on when/how tool can be used, who has access, and audit trails
- ✓ DO: Ensure facial recognition results are not used as definitive identification; require independent corroboration
- ✓ DO: Conduct bias audit before deployment

DON'TS

- ✗ DON'T: Deploy without transparency; disclose use to public and affected communities
- ✗ DON'T: Deploy without testing for accuracy, bias, and real-world impacts in local conditions

Compliance measures:



Restrict circumstances in which biometric data may be used



Civil rights (disparate impact on communities of color)



Vendor contracts (ensure vendor is not misusing city data)

Use Case #4: Benefits Administration & Fraud Detection

OPPORTUNITY:

- Detect fraud
- Recover overpayments
- Improve program integrity
- Operational efficiency and workforce support

REAL-WORLD EXAMPLES & LAWSUITS:

- Several states use AI to flag **suspicious benefit claims** for investigation.
- In 2013, Michigan adopted a fraud detection system, which made 48K accusations against unemployment recipients. **93% of fraud determinations were wrong.**
- California's fraud detection system for unemployment claims was audited in 2022 and found to be **only 46% accurate in flagging fraud.**

CHALLENGES & CRITICISM:

- False positives: Algorithm flags innocent claimants as fraudsters; they lose benefits without due process
- Lack of transparency: Claimants don't know why they were flagged; can't understand/contest the decision
- Disparate impact: If algorithm relies on zip code, credit history, or employment history, it may systematically harm low-income or minority communities

Best Practices: AI for Benefits and Fraud

DOS

- ✓ DO: Provide procedural due process (notice, opportunity to be heard) before denying or clawing back benefits
- ✓ DO: Have a human review AI-flagged claims before final decision
- ✓ DO: Disclose the algorithm's logic to the extent possible
- ✓ DO: Conduct disparate analysis before deploying
- ✓ DO: Allow appeals and provide meaningful way to contest an AI decision.

DON'TS

- ✗ DON'T: Use AI as the final decision-maker on benefit eligibility without human oversight.
- ✗ DON'T: Set and forget. Track error rates, appeal outcomes, demographic patterns.



PART II: GOVERNING PRINCIPLES & COMPLIANCE FRAMEWORK

Common Best Practices Across All Use Cases

1. **Transparency & Disclosure**: Tell the public when AI is being used; what data it uses; what it decides
2. **Human Oversight**: AI informs decisions, but humans make final calls (especially high-stakes: hiring, benefits, enforcement)
3. **Bias & Fairness**: Test for disparate impact on protected classes (race, gender, disability, age)
4. **Data Minimization**: Collect only data you need; delete when no longer necessary
5. **Due Process**: Provide notice and opportunity to contest AI decisions that affect rights/benefits
6. **Vendor Accountability**: Contracts must allocate liability; require vendors to warrant algorithms are not discriminatory
7. **Insurance**: Obtain insurance to cover common risks, including negligence in AI vendor's provision of services, security incidents, and bias and discrimination
8. **Data Security**: Vendors should be subject to robust information security addenda and should be certified to a reputable framework (e.g., SOC II)

Overview of State and Federal AI Regulation

Currently, there is no comprehensive federal AI law. Instead, patchwork of state and local laws in multiple buckets.

- Deepfakes / digital clones
- Companion AI
- Model and data transparency
- Automated decision-making (includes employment plus other high-risk areas)

Across most categories, regulators are focusing on:

- **Transparency** – Disclose when AI is used; explain how it works
- **Accountability** – Vendor and deployer both liable for harmful outcomes
- **Fairness/bias** – Test for disparate impact; audit for bias
- **Data governance** – Minimize data collection; facilitate deletion and access requests
- **Human-in-the-loop** – Especially for high-stakes decisions (hiring, benefits, law enforcement)
- **Citizen redress** – Right to contest AI decision; opportunity to opt-out; meaningful appeal process

Select Proposed AI Bills in Georgia

- **HB 147**: Requires state agencies to maintain and annually report an inventory of their AI use.
- **SB 9**: Expands fraudulent election interference to include publishing (with knowledge of its falsity) “materially deceptive media” within 90 days with intent to deceive electors for certain purposes (including influencing the election).
- **SB 104**: Prohibits state agencies from installing or using any AI system on state equipment that is (i) owned or operated by China (ii) domiciled in or headquartered in China, or (iii) organized under Chinese law.
- **HB 171**: Updates criminal code to prohibit distribution of inappropriate AI-generated content involving children.
- **HB 566**: Prohibits unauthorized use of an individual’s face or voice in digital replicas.
- **SB 167**: Prohibits developers from selling or distributing automated decision-making systems that result in algorithmic discrimination.

Trump AI Executive Order and Federal Preemption Risk

- On December 11, 2025, President Trump signed an Executive Order aiming to limit state AI regulation.
- Rationale: U.S. is in "AI arms race" with China; state regulation slows innovation and U.S. competitiveness.
- Order directs several federal agencies to discourage state AI law promulgation and enforcement through:
 - **Funding conditions** (states with onerous AI laws may not receive BEAD funding)
 - **Risk of litigation** (Attorney General will sue to challenge state AI laws)
 - **Development of framework for a federal-level AI law.**
- The EO does not, on its own, invalidate any state or local AI laws. That would require (1) a court injunction or (2) Congress to enact a federal law preempting those laws.
- Potential shifts?
 - Uncertainty may cause **state legislation to slow.**
 - States may turn to existing consumer protection, unfair competition, and civil rights laws to police AI-related conduct.
 - Risk that companies complying with laws could become the **target of federal enforcement**

AI Compliance Checklist



AI Compliance Checklist

1. Accessibility and Language Requirements:

- a. Provide services in all languages required by law
- b. Ensure ADA compliance
- c. Offer non-digital, human alternatives (phone, in-person)

2. Transparency and User Disclosures:

- a. Clearly disclose that users are interacting with AI
- b. Explain capabilities and limitations upfront
- c. Provide clear escalation path to human staff

3. Data Privacy and Security:

- a. Minimize personal data collection and retention
- b. Evaluate implications of public records laws
- c. Implement appropriate security controls for data in transit and at rest

4. Content and Accuracy Standards:

- a. Implement guardrails on restricted topics (legal, medical, immigration, etc.)
- b. Ground AI responses in official, maintained sources
- c. Update AI content whenever policy or service information changes
- d. Monitor AI for hallucinations

AI Compliance Checklist

5. Risk Assessment and Legal Review:

- a. Conduct impact assessment before launch (privacy, legal, bias)
- b. Document high-risk decisions requiring human review and approval

6. Bias and Fairness Testing:

- a. Audit AI for demographic bias before and after launch
- b. Log and review complaints and error patterns
- c. Document corrective actions and retest regularly

7. Vendor Management:

- a. Ensure responsibilities for compliance steps (e.g., disclosures) are clearly defined by contract
- b. Ensure contractual liability is clearly allocated between the parties
- c. Consider obtaining insurance

8. Governance and Operations:

- a. Adhere to city's AI policy
- b. Ensure vendor complies with AI policy
- c. Designate clear ownership and accountability over vendor management

Questions?



Additional Reading

A Guide to AI Risk Management and Insurance for Modern Companies, Gunderson Dettmer

<https://www.gunder.com/a/web/fQua41PbzbLiY4DNxsgDEz/gd-ai-insurance-white-paper-v7.pdf>



AI in the Workplace: Legal Challenges and Best Practices, presentation slides by Gunderson Dettmer (available upon request to falim@gunder.com).



AI Economics & Readiness

Mathew Garver
Liberty Street Capital, LLC
HiQo Solutions, Inc.

18 February 2026

Mathew Garver

HiQo Solutions, Inc.

- **HiQo Solutions, Inc., CEO**, multinational technology company to Fortune 500 clients worldwide in advanced AI/ML & Edge AI systems
- **Liberty Street Capital, CEO**, private family office growth investment firm for tech enabled companies HiQo, Euromoney, Floor & Décor, Tilray, CGL
- **DARPA**, Advisor on Economic Statecraft
- **AI Council**, Founder, Council for Single-Family Offices
- **Campden Wealth & IPI**, Chairman of Global Summit - NA, UAE, UK,
- **Euromoney** Institutional Investor, Digital Infrastructure, for AI / Edge System
- **Patton Boggs**, Private Capital Group, Emerging Technology for A&D sectors
- **DLA Piper**, Private Capital Group & Infrastructure Steering Committees

EDUCATION

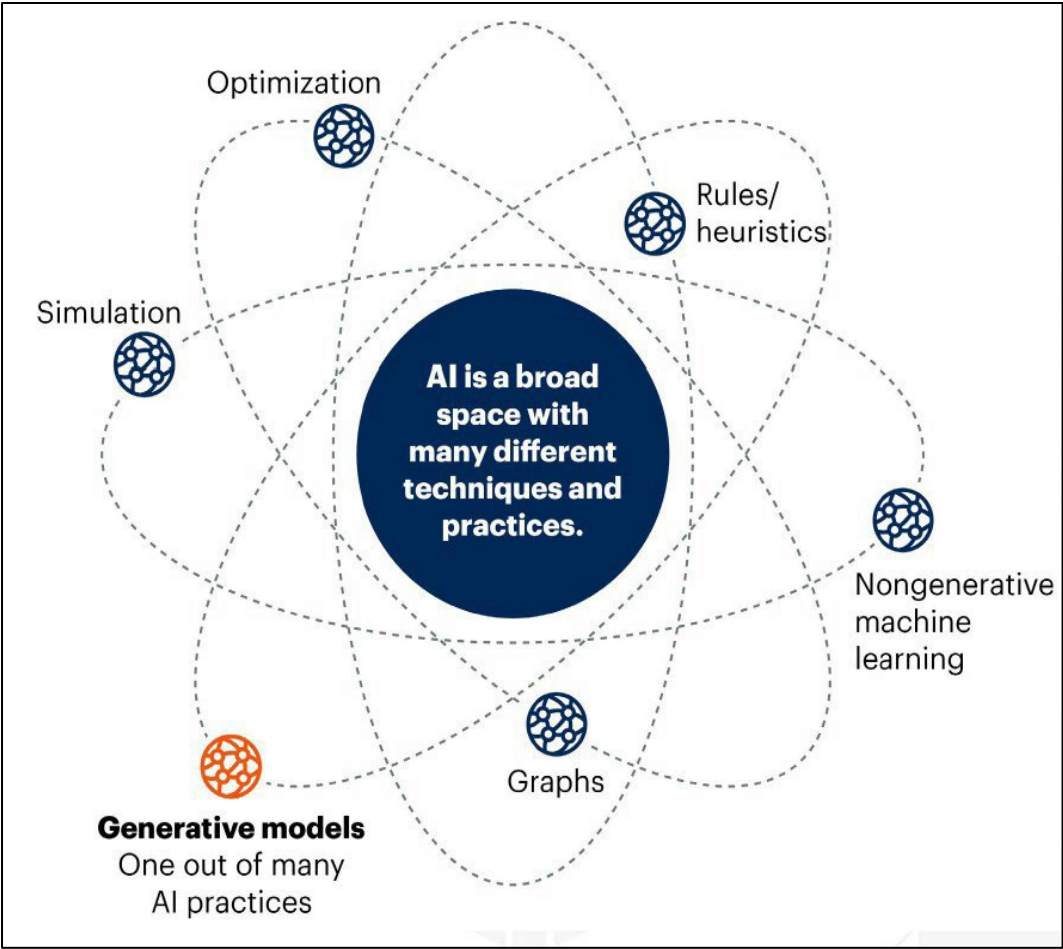
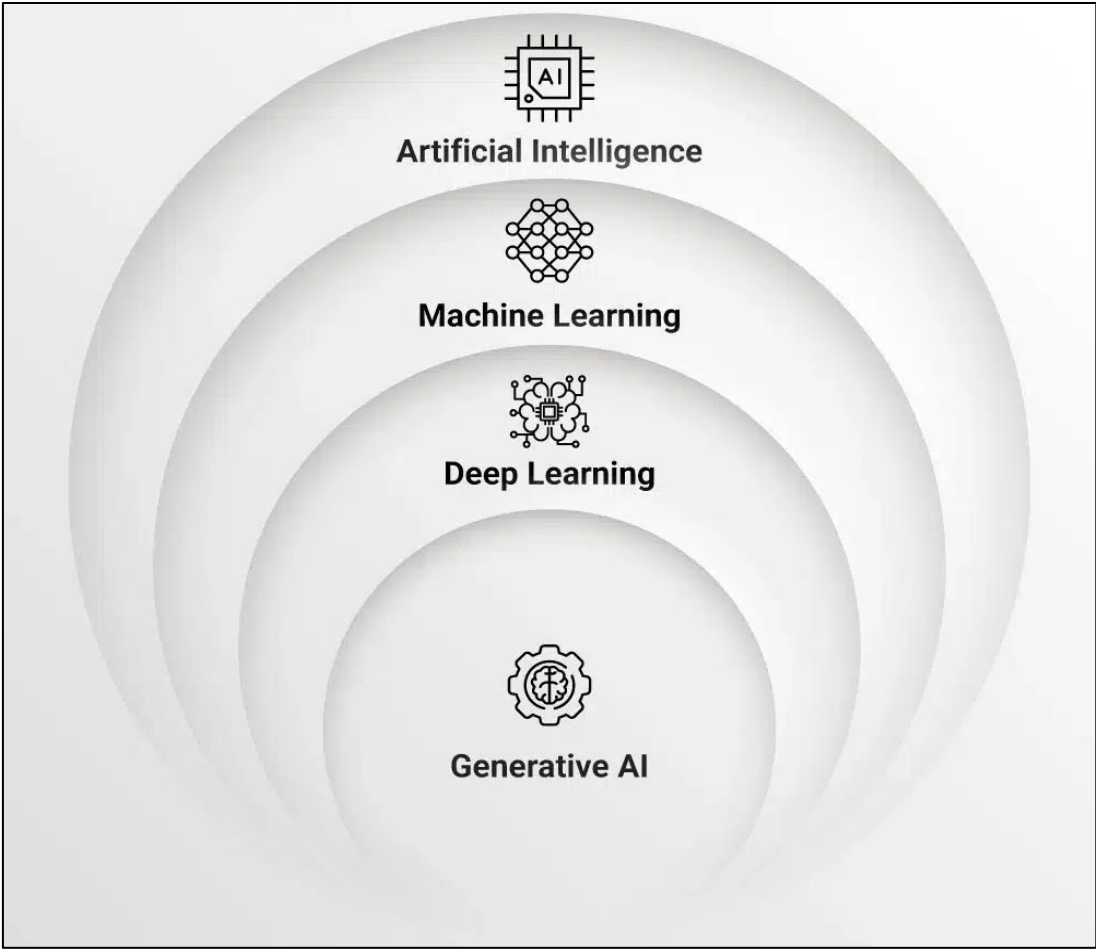
- **Harvard University**, Fellow, International Economics, Weatherhead Center
- **Oxford University** at Saïd Business School; Finance
- **Michigan State University**, Political Economics




Agenda

- The Intelligence Cost Curve
- Software as Ground Zero
- The Jevons Effect: Why Efficiency Creates More Demand
- What AI Readiness Actually Means
- Open vs. Closed System
- Conclusions and Implications for Atlanta


Today's AI Landscape



Commoditizing Models



The Price of Intelligence Is Collapsing




Fastest = 900x / Year


Mid-range = 40x / Year

Slowest = 9x / Year


Marginal Cost for AI Capabilities / Intelligence



OpenAI “o” Series



Google Monthly Tokens Processed




**134x in
18 months**

Exponentially More Capability Per Dollar

Artificial Analysis Intelligence to Price Ratio

Capabilities-to-Cost Curve




Case Study: Software as Ground Zero

If AI can restructure its own industry...

The Industry That Built AI Is the First Restructured by It

12-month Rolling Average of Software Engineering Monthly hires by Top 15 software companies



Source: Zeki, CEO Tom Hurd

Productivity boost — But Where?

Lens 1

Total economic potential of 60-plus organizational use cases¹


Revenue impacts of use cases¹

Cost impacts of use cases

Lens 2

Labor productivity potential across ~2,100 detailed work activities performed by global workforce

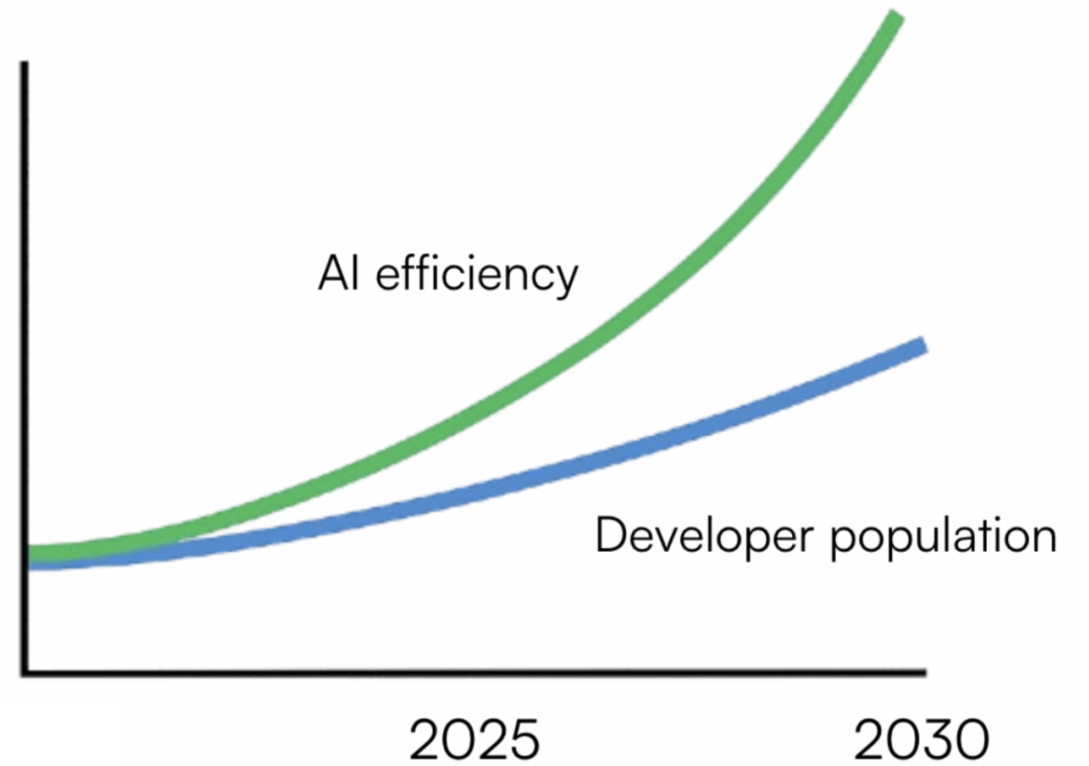

Collapsing Unit Economics



<https://skventures.substack.com/p/societys-technical-debt-and-softwares>

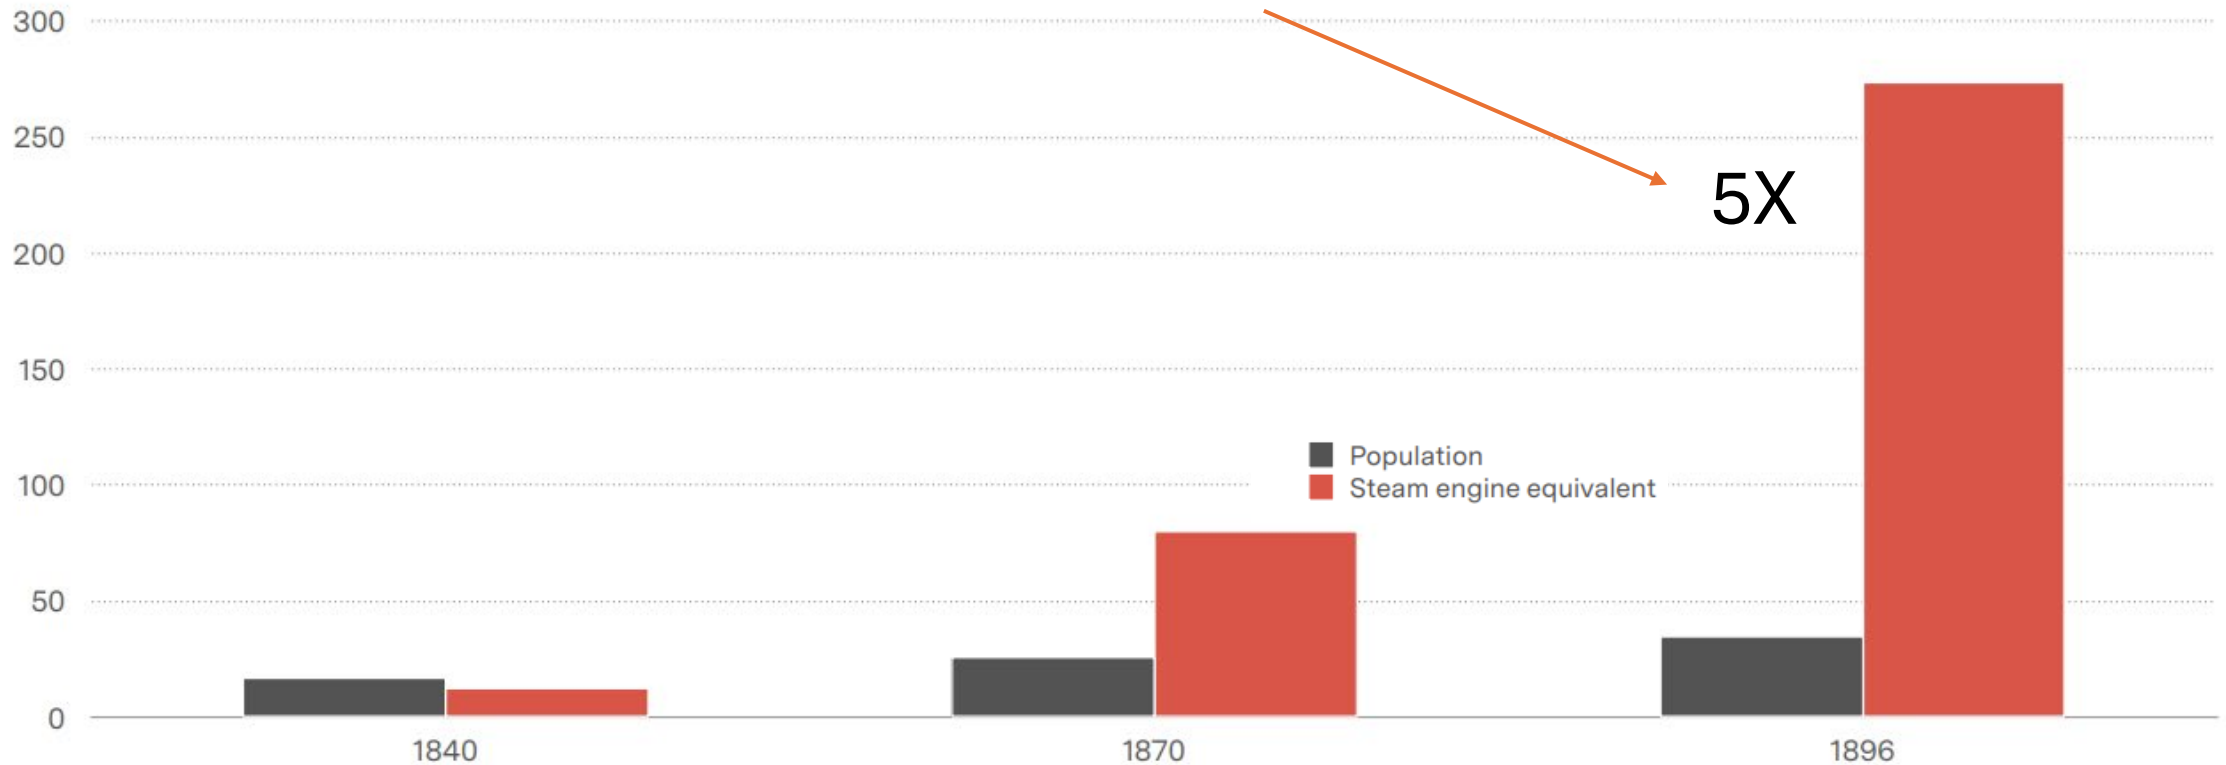
Jevons Paradox

When a resource becomes more efficient to use, total consumption increases — significantly. Jevons observed this with coal in 1865. We are seeing it now with AI.



Jevons Paradox example

Steam Engines gave Britain equivalent labour of
(very roughly) **5X** its total population by 1900




Open Weights and Closed Systems

Open Weights vs Closed (Proprietary) Intelligence

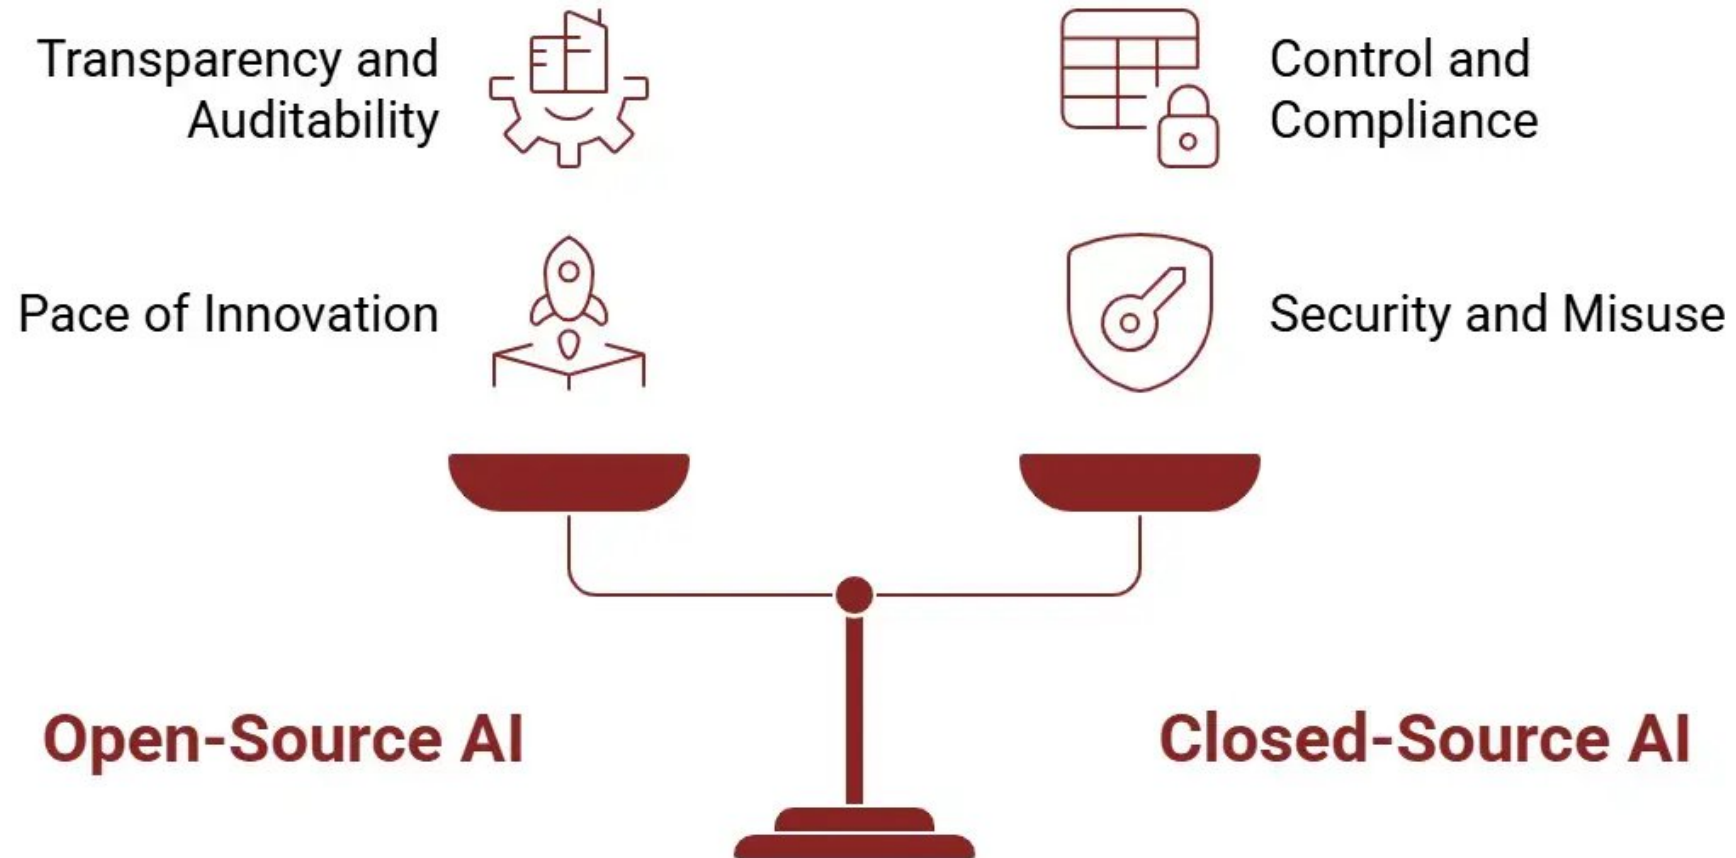
Progress in Open Weights vs. Proprietary Intelligence

Artificial Analysis Intelligence Index v4.0 incorporates 10 evaluations: GDPval-AA, τ^2 -Bench Telecom, Terminal-Bench Hard, SciCode, AA-LCR, AA-Omniscience, IFBench, Humanity's Last Exam, GPQA Diamond, CritPt

■ Open Weights ■ Proprietary



Open vs Closed source, weights



AI Readiness

The gap between AI potential and organizational readiness is a real bottleneck

Paradox: Universal Urgency, Universal Unpreparedness



98% of companies feel increase in urgency



85% of companies believe less than 18 months to deploy AI strategy

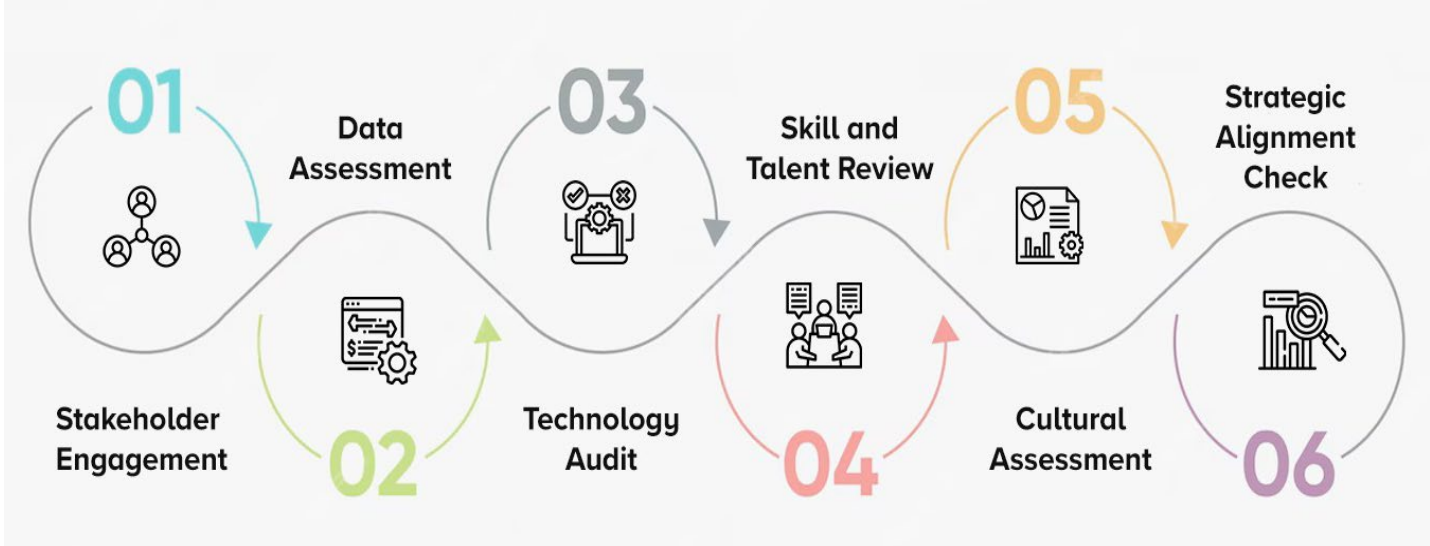
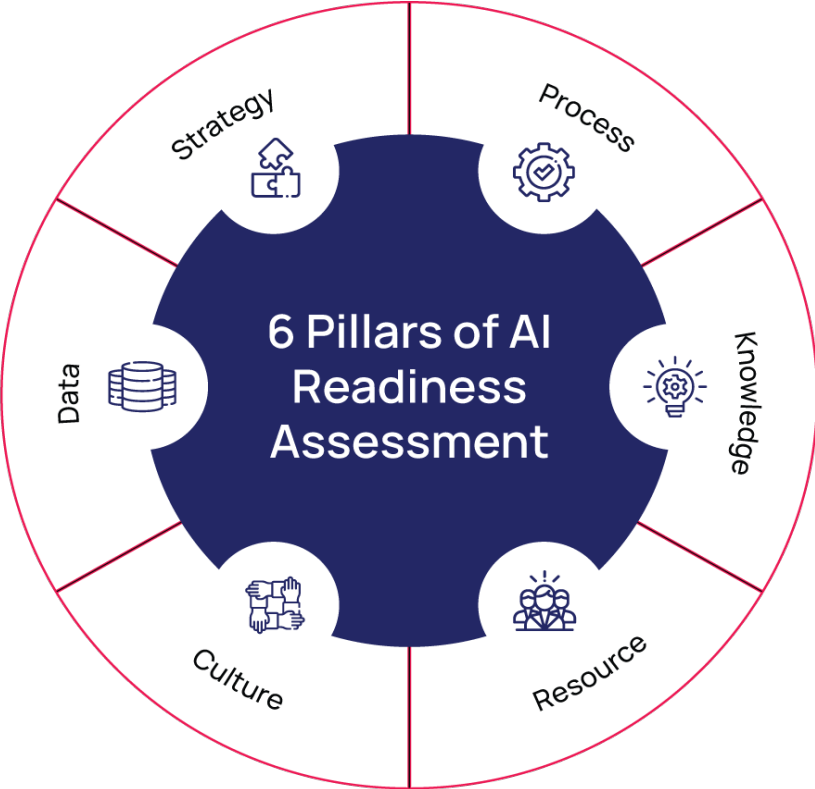


50% of companies have dedicated 10-30% budget to AI

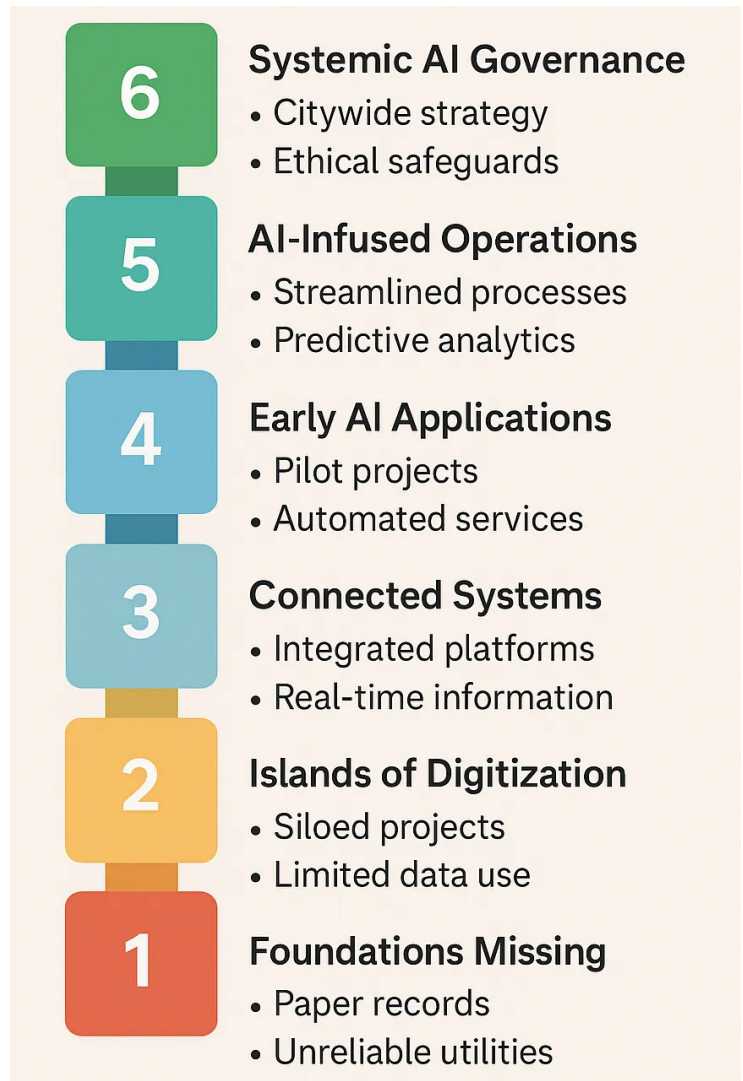


13% of company leadership AI-ready by internal measures

AI Readiness: Alignment Across Pillars



Public Sector AI Readiness



Conclusion: Key Take Aways

1 Intelligence costs are collapsing faster than any input in economic history

- Token prices declining 20-900x/year at the fastest tier. Capabilities doubling every 4-6 months.
- This is not a cycle — it is a structural transition and economic shift.

2 Jevons Paradox: demand for AI will increase — plan for abundance

- As AI gets more efficient, usage explodes — just as steam power multiplied Britain's labor 5X.
- Organizations and governments must prepare for exponentially more AI interactions.


3 Readiness is the bottleneck — and it is an organizational problem, not a technology issue

- Only 13% of company leadership is AI-ready. The gap in government is likely worse.
- Readiness requires aligning data, talent, technology, culture, and governance simultaneously.

Appendices

Why the massive sell off in SaaS Markets?

- Profit Pool in software is expected to shift toward AI Agents
- Illustrative shift in TAM for SaaS and Agents




■ SaaS TAM ■ Agent TAM

Source: Gartner, Goldman Sachs Research
Gartner data published October 10, 2024



**Goldman
Sachs**

Open Weights vs Closed (Proprietary) Intelligence

Open Weight models lag state of the art by roughly 3 months



Open vs Closed economic trajectories







ABOUT EMORY UNIVERSITY'S

CENTER FOR AI LEARNING

Part of Emory's *AI.Humanity* initiative, the Center is our nexus for AI literacy and integration. We bring our students, scholars, and staff to a global community of partners to pave the way to the future, **together**.



46 core tenured and tenure-track AI.Humanity faculty, with a network of 230 faculty with expertise in AI and ML



R1-level expertise in all schools and subject areas, with AI centers and institutes in data science, business, healthcare, nursing, law, liberal arts, and ethics



Member of the U.S. AI Safety Institute Consortium, the cornerstone of the Biden administration's Executive Order on AI, which is setting the national AI research agenda



Our AI training programs engaged more than 2,000 people in 2023; Emory supports more than 15,000 undergraduate and graduate students annually

What are we seeing in business?

How should we think about workforce?

What policies should we implement?

What are we seeing in business?

NEWSLETTERS · CFO DAILY

MIT report: 95% of generative AI pilots at companies are failing



BY **SHERYL ESTRADA**

SENIOR WRITER AND AUTHOR OF CFO DAILY

August 18, 2025 at 6:54 AM EDT



MIT report: 95% of generative AI pilots at companies are failing



BY SHERYL ESTRADA
SENIOR WRITER AND AUTHOR OF CFO DAILY

August 18, 2025 at 6:54 AM EDT



Some of the greatest hits from the field, which I have personally observed:

- **Wrong tool:** teams using the **ChatGPT web interface** to upload and OCR documents. Of course it bombed. You're not supposed to use ChatGPT as OCR. Cloud OCR **APIs** exist, are programmatic, and cut out all the copying, pasting, and uploading.
- **Counterfeit experiments:** "Analyze our A/B test results," except... there was no A/B test. It was a historical slice-and-dice masquerading as causal inference.
- **Strategy leakage:** "Give me a strategy for my industry." By definition, you get the **average** answer your competitors got, too. That's not an edge.
- **No plumbing/governance:** no labels, no feedback loops, no audit trail—we just pump data into a cut and paste LLM prompt, and then when it gives bad advice and poor predictions (based on next best word, by the way), then we blame "hallucinations."

CUSTOMER SERVICE

"Custobots"



The City Website Is Now Part of Service Delivery

The public web is no longer just for humans. If Atlanta does not publish the cleanest machine-readable version of its own services, outside models will become the city's front desk — using stale PDFs and guesswork.

WHAT IS ALREADY HAPPENING

GovInfo launched a public MCP server in Jan 2026 so AI systems can converse with live federal information.

Chrome's WebMCP preview exposes structured tools so agents can act with more precision than raw HTML guessing.

Cloudflare now serves markdown to agents because it's easier and cheaper for AI to process.

Next.js docs describe JSON-LD as structured data that helps AI understand page structure.

Being good at modern web frameworks will matter more than most city leaders realize.

WHAT ATLANTA NEEDS TO DO

Canonical service pages with stable URLs. Clean semantic HTML. JSON-LD structured data on every public service page. Machine-readable forms and FAQs.

This is not a website redesign. It is adding a data layer to what already exists — so that when someone asks an AI agent "how do I get a building permit in Atlanta," the answer comes from the city's own authoritative data — not from a third-party guess.

As AI models get better at writing code, they will build their own MCP connections to any well-structured web presence. Atlanta's web presence either becomes the source — or becomes what AI hallucinates about.

POLICY INPUT Direct the city's web team to implement JSON-LD schema markup across all public service pages. Audit for machine-readability. This is a low-cost action that takes weeks, not months.

Will AI Actually Improve Productivity?

NBER: Access to AI assistance in call centers increases agent productivity by 14%, with the biggest impact on less experienced workers. “We found that workers with access to AI see fairly significant productivity gains, but most of those gains accrue to novice or less able workers,” says Lindsey Raymond, an MIT PhD candidate and coauthor of the new paper. “This may be because the AI model disseminates the potentially tacit knowledge of more able workers and helps new workers move up the experience curve.”

FIGURE 4: EVENT STUDIES, RESOLUTIONS PER HOUR

A. RESOLUTIONS PER HOUR





FIGURE 9: EXPERIENCE CURVES BY DEPLOYMENT COHORT

RESOLUTIONS PER HOUR, BY AGENT TENURE



What Business Has Actually Learned

The strongest early proof is in customer operations – especially contact centers. Here is what the results show:

United Airlines

Real-time agent assist improved response time 15% and cut average handle time 15% in 45 days.

Genesys (own support)

Saved ~5 minutes per interaction. Cut escalations 43%.

Brinks Home

AI coaching cut transfers 73%.

Propel

Reduced after-call work 50% while reviewing 100% of calls instead of a sample.


Klarna

AI handled 2/3 of chats, cut resolution from 11 min to 2 min. Then CEO admitted cost was over-weighted – rehired humans for complex cases.

Bank of America (Erica)

56M contacts/mo. 98% resolved in 44 sec. Over 60% of engagement is now proactive – AI reaches out before customers ask.

The lesson: augmentation first, autonomy second, human fallback always.





Prior Authorization **INCOMPLETE**

📢 Ask if any advanced technology or custom devices will be used.

📢 Ask about bone quality, muscle function, skin, and circulation.

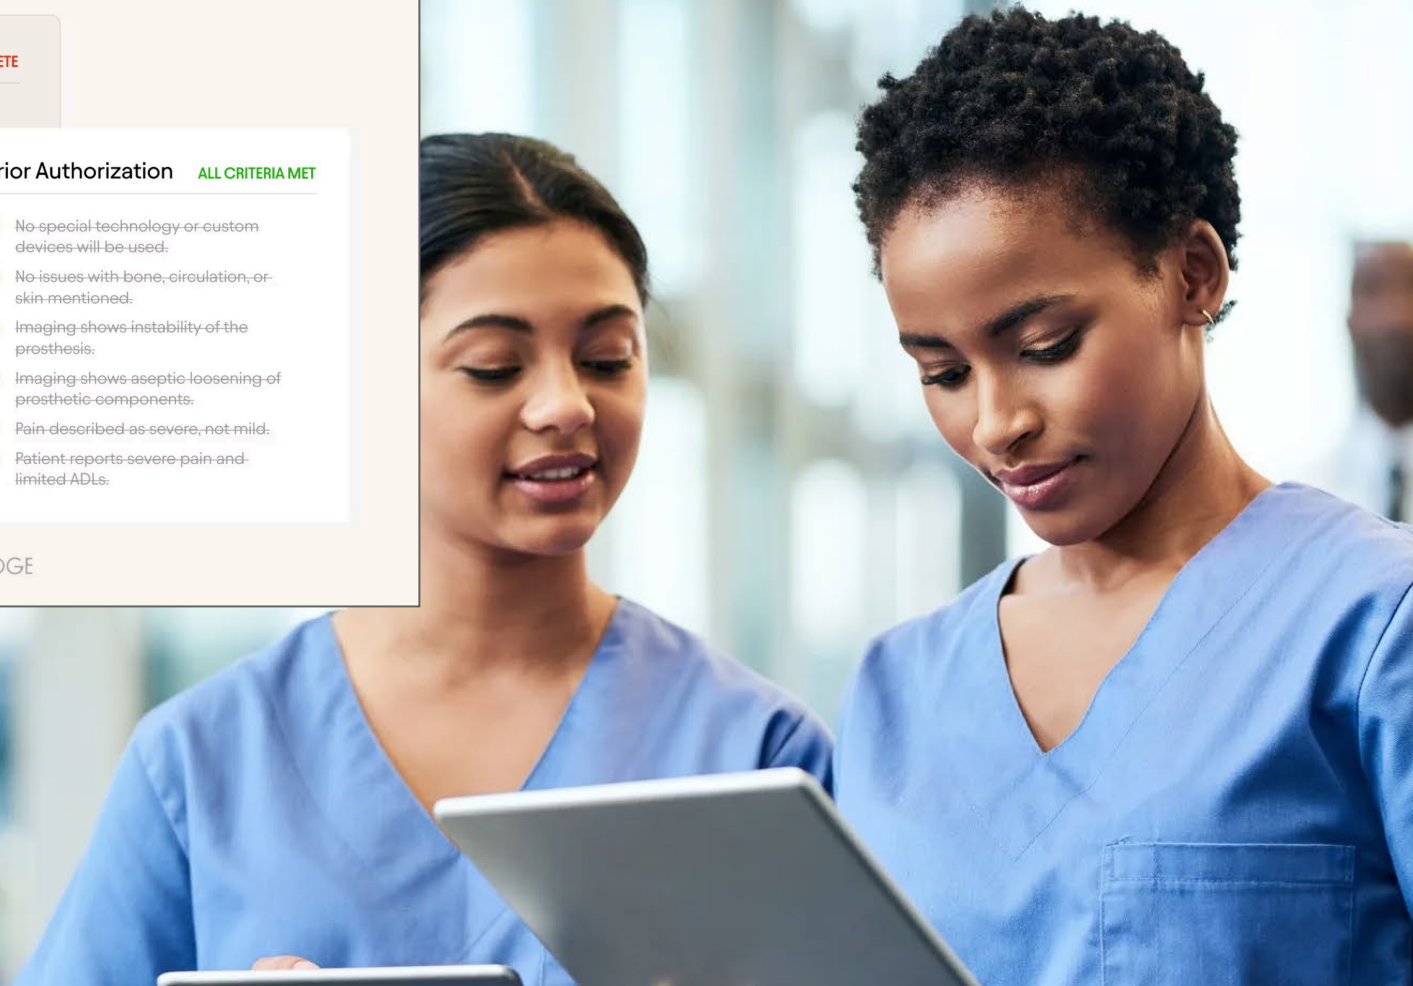
- ✔ Imaging shows instability of the prosthesis.
- ✔ Imaging shows aseptic-loose prosthetic components.
- ✔ Pain described as severe, not mild.
- ✔ Patient reports severe pain and limited ADLs.




Prior Authorization **ALL CRITERIA MET**

- ✔ No special technology or custom devices will be used.
- ✔ No issues with bone, circulation, or skin mentioned.
- ✔ Imaging shows instability of the prosthesis.
- ✔ Imaging shows aseptic loosening of prosthetic components.
- ✔ Pain described as severe, not mild.
- ✔ Patient reports severe pain and limited ADLs.

ABRIDGE







If you liked
HARRY POTTER,
you'll **LOVE** these
books by Random House



 **RANDOM HOUSE**





2024



workday.







@JoeSutherlandInsta

What are we seeing in workforce?

THE HILL

Almost 100M jobs could be lost to AI, automation: Senate report

BY MAX REGO - 10/06/25 4:06 PM ET



We Are Still Underreacting on AI

This is not just a technology issue, it's a fundamental change to our society—and we remain dangerously underprepared.

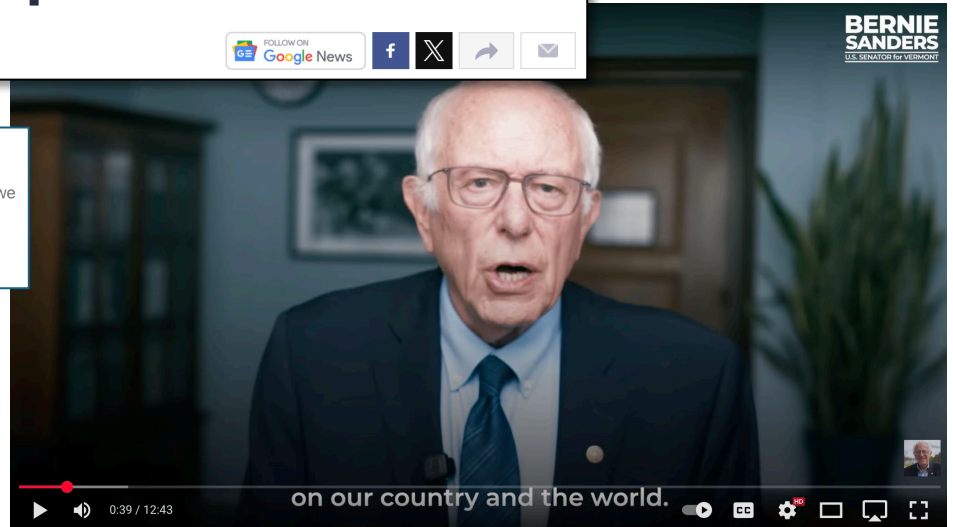


PETE BUTTIGIEG
JUN 24, 2025

How your kids will make money

In a world of AI

ROBERT REICH
AUG 19, 2025




What does ChatGPT Predict?

To understand how artificial labor could impact workers and their livelihoods, HELP Minority Staff asked ChatGPT about job displacement. Specifically, HELP Minority Staff asked ChatGPT to analyze job descriptions catalogued by the federal government for the entire U.S. economy and predict tasks that could be performed by AI and automation.

Table 2: ChatGPT's prediction of the 20 occupations with the largest job losses, based on tasks performed, ranked by the total number of jobs replaced by artificial labor over 10 years.

Job Description	Number of Jobs Replaced	% of Jobs Replaced
Fast Food and Counter Workers	3,310,372	89%
Customer Service Representatives	2,462,167	83%
Laborers and Freight, Stock, and Material Movers, Hand	2,441,400	81%
Retail Salespersons	2,372,957	62%
Stockers and Order Fillers	2,170,227	76%
Cashiers	1,975,457	59%
Office Clerks, General	1,736,306	66%
General and Operations Managers	1,721,971	47%
Secretaries and Administrative Assistants, Except Legal, Medical, and Executive	1,580,658	80%
Home Health Aides	1,574,601	40%
Janitors and Cleaners, Except Maids and Housekeeping Cleaners	1,473,697	61%
Registered Nurses	1,309,563	40%
Bookkeeping, Accounting, and Auditing Clerks	1,257,995	76%
Waiters and Waitresses	1,208,682	53%
Heavy and Tractor-Trailer Truck Drivers	1,031,940	47%
Accountants and Auditors	992,214	64%
Personal Care Aides	973,098	25%
Team Assemblers	924,408	61%
Software Developers	920,265	54%
Teaching Assistants, Preschool, Elementary, Middle and Secondary School, Except Special Education	896,576	65%


Employee headcount among software developers, by age



Note: Indexed to 100 at October 2022

Source: Brynjolfsson, Chandar and Chen

Employee headcount among software developers, by age




Note: Indexed to 100 at October 2022

Source: Brynjolfsson, Chandar and Chen

What else happened in 2022?

Aggressive Rate Hikes
Russia Invades Ukraine
Inflation Whiplash
COVID/RTO
Huge Crypto Winter
Euro Devaluation
Stocks Hit
Trade War Reignites
and more.....(!)

Employee headcount among software developers, by age

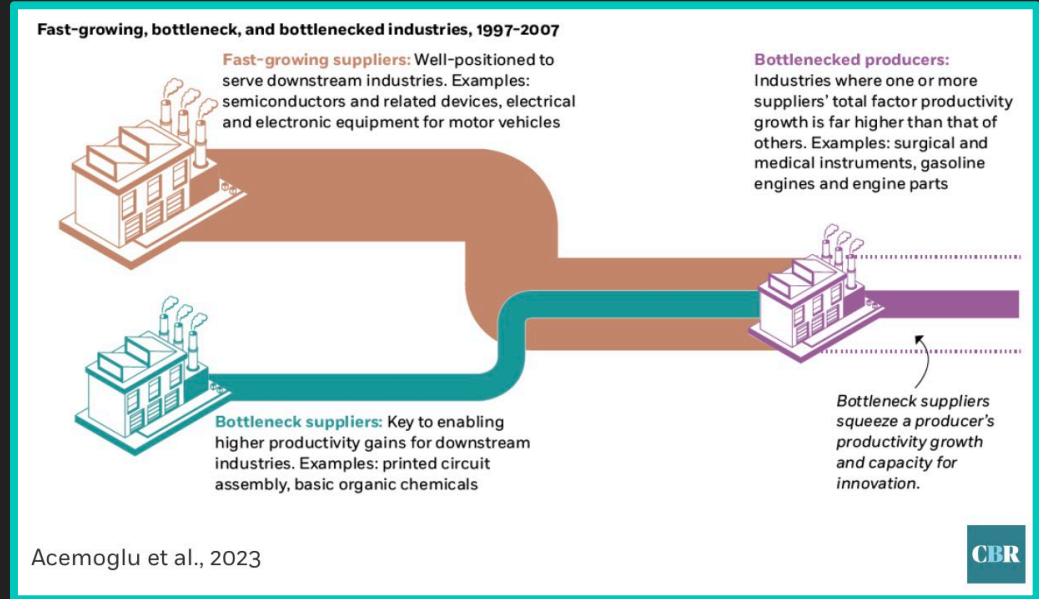


Note: Indexed to 100 at October 2022

Source: Brynjolfsson, Chandar and Chen

7 Reasons We See Low Growth:

1. Technological Progress Slowed
2. *Diffusion Highly Constrained*
3. Zombie Firms/Misallocation
4. Financial Headwinds / No R&D
5. Overregulation v. Concentration







AI + YOU



EMORY
UNIVERSITY

Center for AI Learning



ROWEN™



GEORGIA CHAMBER



AI + YOU

How will AI affect my life? My kids' lives?

Will AI replace our jobs?

Is my privacy at risk?

Is AI sentient?

How will AI change our education system?

What are we doing to regulate AI?

How can I access educational resources?

70%

of college
graduates wish AI
had been
integrated into
their courses

10%

of higher
education
presidents and
provosts have an
AI strategy

AI + YOU

What You Can Do

1. **Clear Strategy:** Atlanta's workforce competitiveness in AI will come down to strong, accountable leadership in education.
2. **AI Literacy:** Every teacher/faculty should have access to training programs and AI tools without putting down a credit card.
3. **Workforce Development** is a must: not just STEM course grants. We're building new programs cost-effectively through partnerships that can be used broadly (e.g. Rowen). Online!
4. **Access:** Wage increases, worker retention, and innovation will be broad-based only with broad access to broadband and AI tools; siloed access will limit productivity and growth.
5. **Less Town-and-Gown, More Experiential Learning:** Investment in public-private innovation hubs will produce talent and tools that solve the state's biggest problems in healthcare, human services, manufacturing, agriculture.

Higher Ed as the City's AI Workforce Engine

Atlanta has something most cities don't: a concentration of research universities, healthcare systems, and a community AI literacy program that's already at scale.

WHAT EXISTS TODAY

Emory Center for AI Learning: \$300M research portfolio. Annual AI conference, 400+ attendees. NIST/federal policy work through CAISI.

AI + You Georgia: Free community AI literacy tour. 22+ cities across Georgia. Sold out in every market. Events held at local technical colleges and chambers of commerce. Attendees get 3-month AI platform access and can earn an Emory AI certification.

Emory Healthcare: One of the largest health systems in the Southeast. AI deployment in clinical settings creates a direct knowledge pipeline for how AI performs in high-stakes, regulated environments.

Georgia Tech, Morehouse, Spelman, Clark Atlanta, GA State: All within 10 miles of City Hall.

THE MODEL FOR ATLANTA

1. Workforce development pipeline: AI + You Georgia's model – community events, free tools, certification – should be formalized via a city AI literacy partner. Train city employees, small-business owners, and residents.
2. Research-to-practice bridge: Emory and Georgia Tech evaluate city AI deployments before and after launch. Independent assessment, not vendor claims.
3. Healthcare as proving ground: Emory Healthcare's experience governing AI in clinical settings – bias testing, human oversight, outcome monitoring – is directly transferable to city services.
4. Talent supply: Co-op and fellowship programs that embed graduate students in city departments during early deployments. The city gets capacity. Students get real-world experience.

POLICY INPUT Formalize an Atlanta AI Partnership with Emory, Georgia Tech, and the TCSG system. Include independent deployment evaluation, workforce training, and a talent pipeline for city departments.

What policies are we interested in?





Key Partners

A key part of CPSAI's mission is be connective tissue

Procurement Framework!



Health & Human Services Agencies

NIST and CAISI as Operating Rules, Not Slogans

Formerly the U.S. AI Safety Institute (AISI). Renamed CAISI in June 2025 under Commerce Secretary Lutnick. Same core work – voluntary testing, guidelines, industry collaboration – now with a sharper focus on competitiveness and national security. These are voluntary guidelines, not mandatory rules. But they are the best framework available.

AI RMF 1.0: Govern → Map → Measure → Manage

Voluntary framework for trustworthy AI. The city's governance should align to this four-function structure.

CAISI Agent Standards Initiative (Feb 2026)

Focused on trusted, interoperable, secure AI agent systems. Listening sessions for healthcare, finance, education starting April.

Post-Deployment Monitoring (AI 800-4)

Six things to watch after launch: does it still work, are operations stable, how do people use it, is it secure, is it compliant, what are the broader effects.

THE MOST USEFUL MOVE FOR ATLANTA Propose a one-page Atlanta AI Project Card for every deployment. Seven fields:

1. What problem it helps with
2. Who is affected (residents, businesses, staff)
3. What source material is authoritative
4. What happens when it is wrong
5. What gets watched monthly
6. Who can stop it
7. How a resident or business can challenge the outcome

This is the NIST framework made operational – not as a poster, but as a living document attached to every deployment.

Read. Guide. Act.

Every successful AI deployment in service operations follows three moves. This becomes the spine for everything else in this deck.

READ

The system digests calls, emails, PDFs, permit packets, ordinances, contracts, and case notes. It turns unstructured information into something a worker – or another system – can use.

Transcribe a 311 call. Parse a 40-page permit application. Summarize a contract clause.

GUIDE

It helps a worker in the moment: what rule applies, what question comes next, what language to use, when to escalate. The human stays in the loop. The AI makes them faster.

Surface the relevant code section during a live inspection. Suggest the right routing for a cross-department 311 request. Coach a new hire through a complex permit question.

ACT

It does a few bounded things across systems: opens a case, drafts a letter, routes a request, updates a status. Always within defined guardrails. Never makes a final decision on enforcement, benefits, or rights.

Draft a deficiency notice. Auto-route a service request. Update a permit status. Generate a case summary after a call.



Computing

Neural Network Learns to Identify Criminals by Their Faces

The effort aimed at identifying criminals from their mugshots raises serious ethical issues about how we should use artificial intelligence.

[arXiv.org > cs > arXiv:1611.04135](https://arxiv.org/abs/1611.04135)



Computer Science > Computer Vision and Pattern Recognition

Automated Inference on Criminality using Face Images

Xiaolin Wu, Xi Zhang

(Submitted on 13 Nov 2016 (v1), last revised 21 Nov 2016 (this version, v2))

MIT
Technology
Review



(a) Three samples in criminal ID photo set S_c .



(b) Three samples in non-criminal ID photo set S_n .

Computing

Neural Network Learns to Identify Criminals by Their Faces

The effort aimed at identifying criminals from their mugshots raises serious ethical issues about how we should use artificial intelligence.

[arXiv.org > cs > arXiv:1611.04135](https://arxiv.org/abs/1611.04135)




Computer Science > Computer Vision and Pattern Recognition

Automated Inference on Criminality using Face Images

Xiaolin Wu, Xi Zhang

(Submitted on 13 Nov 2016 (v1), last revised 21 Nov 2016 (this version, v2))

MIT
Technology
Review



Cesare Lombroso, c. 1876

Computing

Neural Network Learns to Identify Criminals by Their Faces

The effort aimed at identifying criminals from their mugshots raises serious ethical issues about how we should use artificial intelligence.

[arXiv.org > cs > arXiv:1611.04135](https://arxiv.org/abs/1611.04135)




Computer Science > Computer Vision and Pattern Recognition

Automated Inference on Criminality using Face Images

Xiaolin Wu, Xi Zhang

(Submitted on 13 Nov 2016 (v1), last revised 21 Nov 2016 (this version, v2))



MIT
Technology
Review




Cesare Lombroso, c. 1876

This is complete BS!!!!

Source: BS Course

  @JoeSutherlandInsta





@JoeSutherlandInsta

What Nobody's Seeing — and What's Finally Possible

WHAT CITIES AREN'T SEEING YET

AI agents will create machine-to-machine demand at a scale current service rules were never written for: status checks, applications, vendor questions, maybe even public comments.

Applicants will use AI to optimize submissions against city rules.

External documents can carry malicious instructions into internal tools (prompt injection via uploaded PDFs).

Deepfake testimony at public hearings. No city has a verification policy.

If the city doesn't publish clean authoritative information, third-party models become the unofficial front desk.

WHAT CITIES CAN FINALLY DO

Cities have spent 20 years digitizing forms and websites and still failed at one plain thing: helping people understand government.

AI may finally let them:

- Create one front door across siloed departments
- Keep information current in multiple languages
- Review every interaction instead of 1%
- Preserve institutional memory despite turnover
- Turn dense code and policy into live guidance in the moment of service

Government has digitized submission. It may now be able to digitize understanding.

POLICY INPUT Draft policy positions on AI-generated public testimony, machine-to-machine service demand, and equity testing requirements now. Don't wait for the incident to write the rule.

AI and Public Data



Charlotte S. Alexander

Professor of Law and Ethics

Georgia Tech Scheller College of Business

Law, Data & Design Lab


Introduction



Law, Data & Design Lab

Developing Innovative Solutions to Justice Problems

The Law, Data, and Design Lab is an interdisciplinary group of problem-solvers who are working to increase fairness, efficiency, transparency, and access to justice in the civil and criminal legal systems in the United States and around the world.



Justice Problems

THE WALL STREET JOURNAL.

Clogged Courts Worsen U.S. Immigration Crisis

By [Alicia A. Caldwell](#)

Sept. 1, 2023 5:30 am ET

The Atlanta Journal-Constitution

Fulton continues hiring to reduce court case backlog

By Jim Gaines

Nov 25, 2023

The Washington Post

Researchers have discovered a new and surprising racial bias in the criminal justice system




By Jeff Guo

February 24, 2016 at 8:48 a.m. EST

Data Problems

“[C]ourts—unlike businesses—**do not usually create data** about their operations and analyze such data to improve the courts’ practices.”



UCLA ANDERSON
REVIEW

Research Brief

A Major Medical Center Gets a Cheaper, Fairer Way to Assign Doctors

Model tells schedulers which anesthesiologists should be on call or on-site at specific times

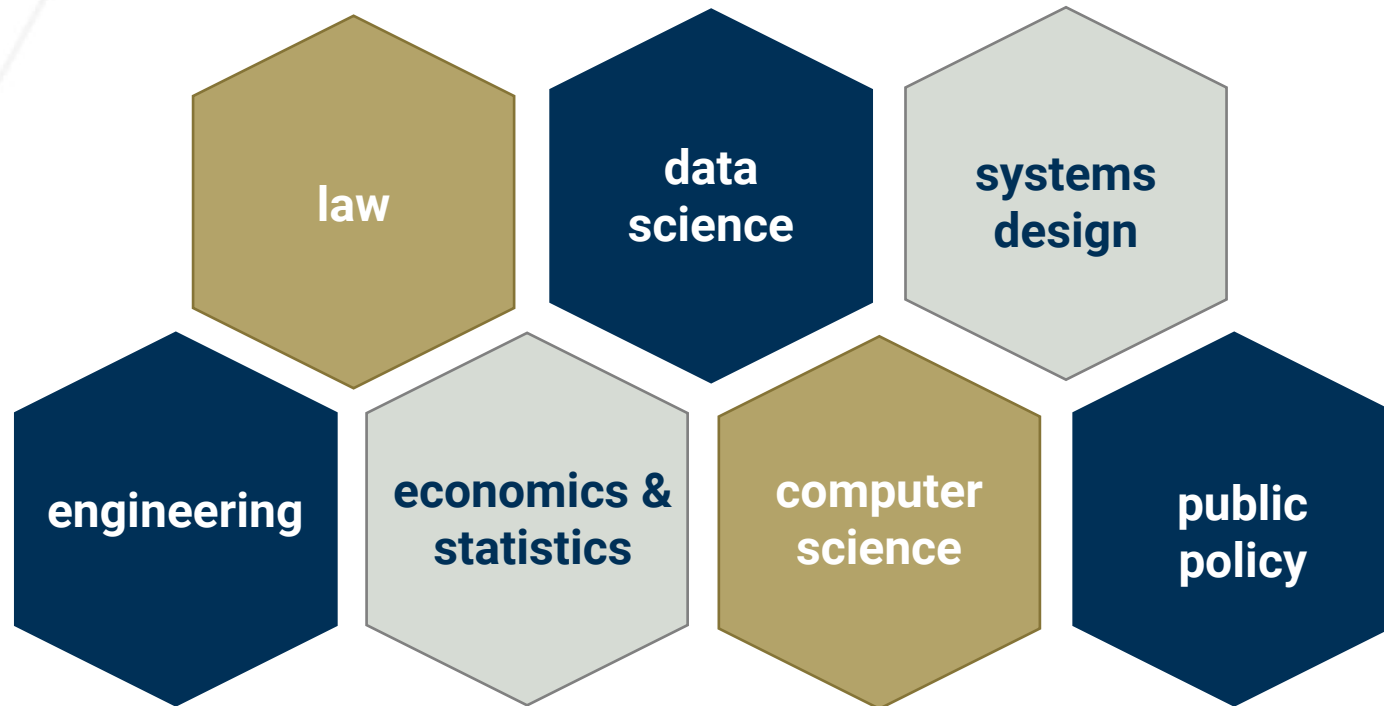


How UPS' ORION Algorithm Transformed Its Route Optimization

Obi Tabansi | 26 July 2025 | 7 min read



Access to Data = Access to Justice



INSIGHTS
POLICY FORUM

LEGAL RESEARCH

How to build a more open justice system

Court records are unstructured and costly to access—here's how to fix it

By Adam R. Pahn¹, David L. Schwartz¹, Sarah Sanga¹, Zachary D. Clopton¹, Peter DiCola¹, Rachel Davis Mersky², Charlotte S. Alexander³, Kristian J. Hammond⁴, Luis A. Nunes Amaral¹

Modern governments gather information across an extraordinary range of activities and use this information to direct policy. Whether a central bank monitoring inflation or a health agency monitoring disease, these entities typically publicly disclose the information gathered so that their actions can be reviewed and evaluated by others. But in many respects, the justice system is a glaring exception. In the United States, a range of technical and financial obstacles blocks large-scale access to public court records—all but foreclosing their use to direct policy. Yet a growing body of empirical legal research demonstrates that systematic analyses of court records could improve legal practice and the administration of justice. And although much of the legal community resists quantitative approaches to law, we believe that even the skeptics will be receptive to quantitative feedback—so long as it is straightforward, apolitical, and incontrovertible. We offer an example of this kind of feedback as well as a collaborative research agenda to dismantle access barriers to court records and enable the public to analyze them.

Although court records in the United States sit in the public domain, federal courts charge \$0.10 per printed page to view any record online (1). Accessing a single case might cost \$10 or more. Accessing all cases from a given year would cost millions of dollars (2). To be sure, the federal judiciary releases in-house studies that use federal court records, as well as a database of basic information

154 | 10 JULY 2020 • VOL. 8/9 ISSUE 6500

Published by AAAS

science.org SCIENCE



Example Project: Case Linkages

Data & Results

- 1,000 criminal defendant names from Fulton County courts
- AI-enabled name matching in civil cases
- ~2,400 total associated cases

Fulton County Magistrate, State, and Superior Court Record Search

Welcome, Charlotte

Smart Search *Required


Search Criteria ?

* Enter a Record Number or Name in Last, First Middle Suffix Format

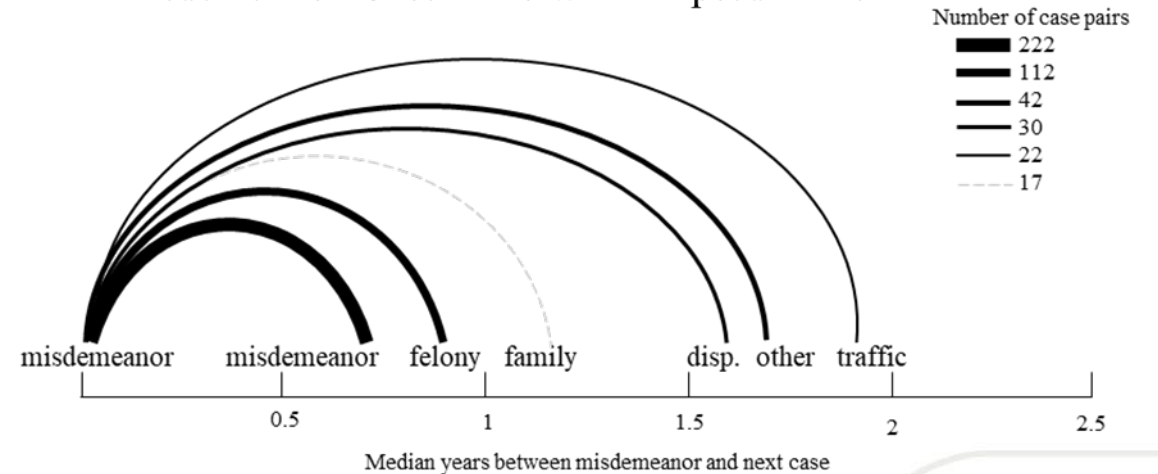
Clear Submit

[Advanced Filtering Options](#)

Felony Case Pairs with Elapsed Time



Misdemeanor Case Pairs with Elapsed Time

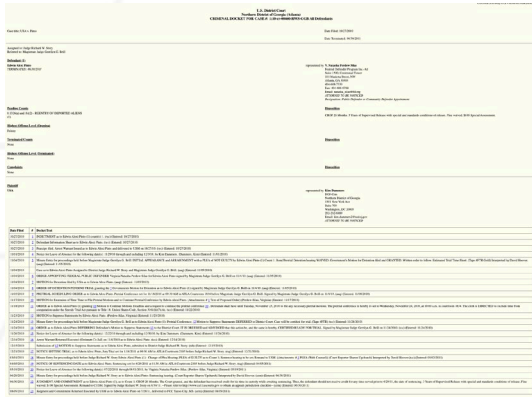


A vintage white car with a Georgia Tech logo on the side, parked on a lawn in front of a brick building. The car is a classic open-top model with spoked wheels and a spare tire mounted on the side. The background shows a large, multi-story brick building with many windows, partially obscured by trees. The entire scene is overlaid with a semi-transparent light green filter.

Example Project: Court Delays

Data

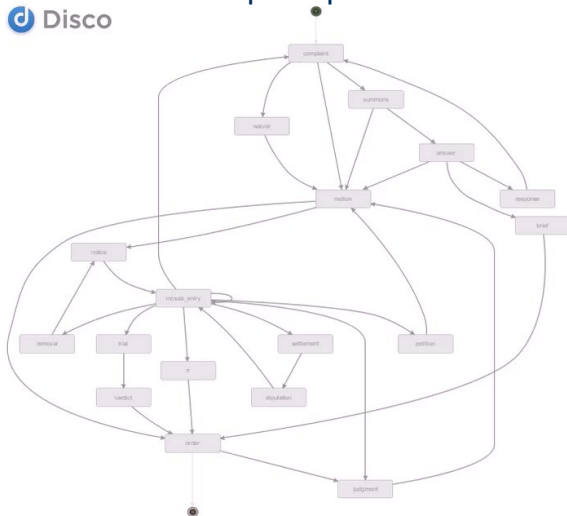
Original data:
Docket sheet text



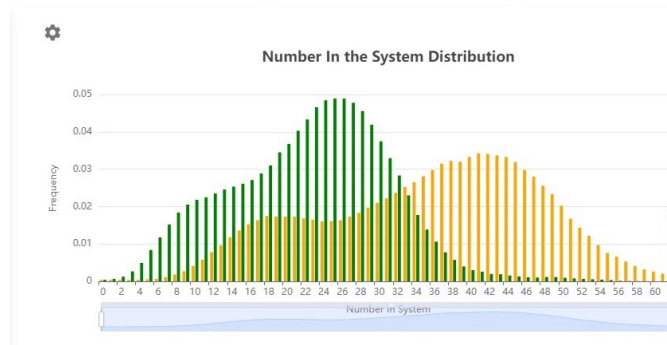
Event log

Case ID	Timestamp	Activities	Server ID	Activity Attributes	Case Attributes
1:02-cv-00430	2002-02-01	complaint	SJ000169	opening	Employment
1:02-cv-00430	2002-02-19	summons	SJ000169		Employment
1:02-cv-00430	2002-03-08	motion	SJ000169	motion to dismiss	Employment
1:02-cv-00430	2002-03-13	order	SJ000169	hearing	Employment
1:02-cv-00430	2002-06-26	answer	SJ000169		Employment
1:02-cv-00430	2002-08-14	order	SJ000169	scheduling	Employment
1:02-cv-00430	2002-12-05	motion	SJ000169	time extension	Employment
1:02-cv-00430	2003-01-16	minute entry	SJ000169	unopposed	Employment
1:02-cv-00430	2003-01-16	order	SJ000169	rescheduling	Employment
1:02-cv-00430	2003-08-22	settlement	SJ000169	dispositive	Employment


Process mining:
Map the process



Queue mining: Analyze,
diagnose, predict delays



Interventions:
Simulated and real



PREPARED FOR

City of Atlanta
AI Commission

AI · GOVERNANCE · PUBLIC INFRASTRUCTURE

Governance in the AI Era

How Atlanta can approach governance infrastructure
that protects residents and strengthens local control

CONTEXT

Why this conversation now?

We are responding to signals, not settled outcomes.

- Every city is being asked to have an opinion on AI, often without a clear definition of what that means.
- The pressure to adopt AI is real, but long-term impacts on public institutions are still unfolding.
- Rapid technical change is a constant, but institutions move on budget cycles, governance, and administrative changes.
- The goal is never to predict the future. Instead, cities can make intentional decisions now that hold up against time and change.

*The most important question isn't whether to use AI but **how we govern the data that AI utilizes and the decisions that it supports.***

FRAMING

What do we mean by "AI"?

There is no consensus on where current trends lead, how fast they move, or its institutional impact.

Automation

Machine Learning

Generative Tools

Workflow Acceleration

The Acceleration View

- Rapid iteration cycles
- Increasing model capability
- AI as a force multiplier
- Personal productivity gains

The Skeptical View

- Persistent limitations
- Models still produce errors
- Hype exceeds readiness
- Brittle and uneven

Is AI ever going to stop changing?

No. Capabilities, risks, and tools evolve continuously. Cities must build governance that absorbs change instead of trying to respond to it retroactively.

INSTITUTIONAL REASONING

Choosing how to think about technology

Public institutions are not just choosing technologies.

AI adoption is ultimately about institutional reasoning.

*Cities need **thinking tools**, not just software.*



What is the state of the technology?

Separating capability from hype



What decisions should it influence?

Scope before deployment



Where must human judgment remain primary?

Boundaries that protect public trust



What risks emerge when systems scale?

Anticipating downstream effects

Before adopting tools, institutions need clarity about what they are dealing with and a framework for making decisions that outlast any single technology cycle.

Governance determines whether AI strengthens or strains institutions.



Governance as structure





Acceleration without guardrails increases systemic risk.

THE RISK

- › Data misuse
- › Malicious interfaces
- › Reduced oversight
- › Erosion of public trust

Without good governance, AI can easily be **extractive**.

GOVERNANCE STRUCTURES

-  Clear participation agreements
-  Defined decision authority
-  Role-based access
-  Auditability & transparency

With good governance, AI can be held **accountable**.

Governance is the stabilizing force. Without it, technological acceleration increases systemic risk across every public system.

GOVERNANCE

What does “AI governance” mean for cities?

Much more than the AI systems alone

- **Data governance:** standards for data quality, access, classification, ownership, and lifecycle management across municipal systems.
- **Ethics governance:** establishing principles and review boards for fairness, bias, transparency, and accountability in AI use.
- **Privacy governance:** how resident data are collected, stored, shared, and protected, especially when feeding AI systems.
- **Regulatory/compliance governance:** aligning AI use with federal, state, and local legal requirements.

Some of these governance domains already exist inside city government, but rapid advancements in AI change the speed, scale, and visibility demands of each.

DATA GOVERNANCE

AI and public data

Data governance is the upstream policy decision.

Cities collect this data every day:



HIPAA-regulated
health data



Social determinant
data



Utility
data



Benefits
enrollment

Governance questions have to start before AI enters the picture.

Example: Summer EBT. When a federal nutrition program requires automatic enrollment utilizing records covered by FERPA, an administrative requirement becomes a governance problem.

PRACTICAL EXAMPLES

Asemio work in practice

*Governance-first design.
Responsible data integration.
Infrastructure that absorbs change.*



Eviction Early Warning

Combining court, utility, and social services data to identify at-risk households before a filing occurs.



Summer EBT & Tribal Sovereignty

Managing data sharing across sovereign and state systems when federal and state funding diverged.



Tulsa: Health + Community Data

Linking clinical and social determinant data via the mayor's health council with local control preserved.



City Water Utility Data

Turning routine operational city data into a strategic community needs assessment tool.

WRAPPING UP

Takeaways for Atlanta

Local communities can maintain control and governance while still collaborating for improved insight, even in a world with AI.

Three Questions to Take Back

1. Are we being intentional about how we govern the data we already collect?
2. Do we have insight into which AI tools are already in our workflows?
3. Do we have a clear path to the strategic use of data assets and AI tools?

About Asemio

We are the technology, data, and consulting firm for mission-based nonprofit, government, and philanthropic organizations that are addressing today's social problems and underserved communities. With a special focus in integrated data system (IDS) development, we equip social impact teams with the data they need to do good.

