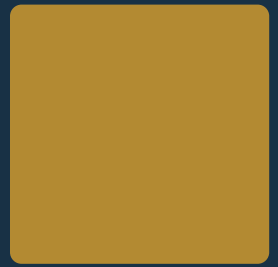
Making the most of AI for government

Awareness, use cases, and risk controls for local government operations





OASBO Summer Conference 2026

About me...

Stuart Long

Chief Information Officer for Cascade Technology Alliance (Northwest Regional ESD & Multnomah ESD)

28 year career in education, technology, & technology leadership

Extensive work on regional networking, cybersecurity, & AI

MS in Applied Information Management, University of Oregon

Certified Information Systems Security Professional (CISSP)

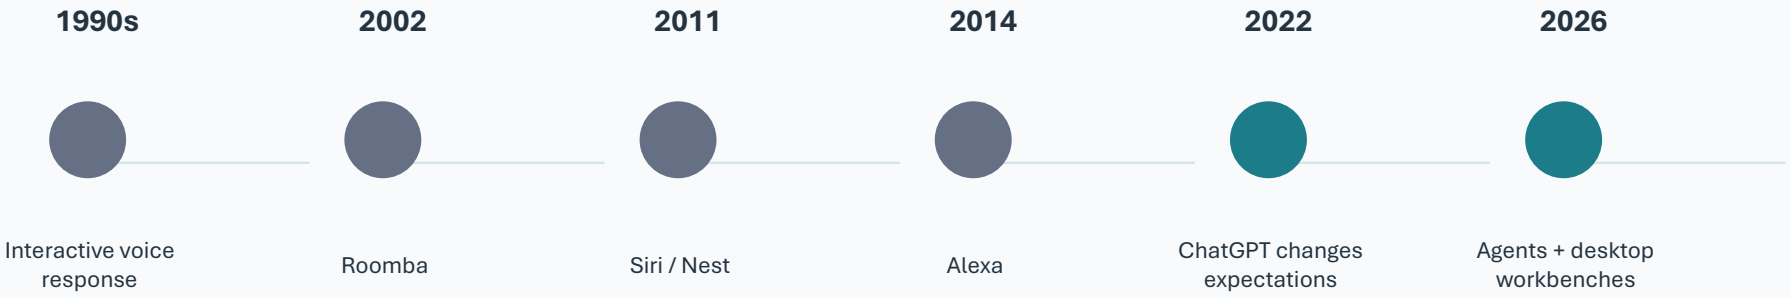
ORIENTATION

Today’s agenda

- Understand what modern AI is — and what it is not
- Know where AI is useful in local government operations
- Awareness & tools to manage AI risk
- Recognize the new 2026 shift: from chatbot to workbench / agent
- Leave with a safe adoption playbook for faster work, but not automated judgment

BASICS

AI has been around for years



The dawn of “AI” in 2022 was the debut of a natural-language interface that felt relatable.

And in 2022...

- OpenAI introduces ChatGPT
  - GPT - Generative Pre-trained Transformer
- Focused on natural language interface
- Limitless topics
- “Magical” ability for natural human conversation

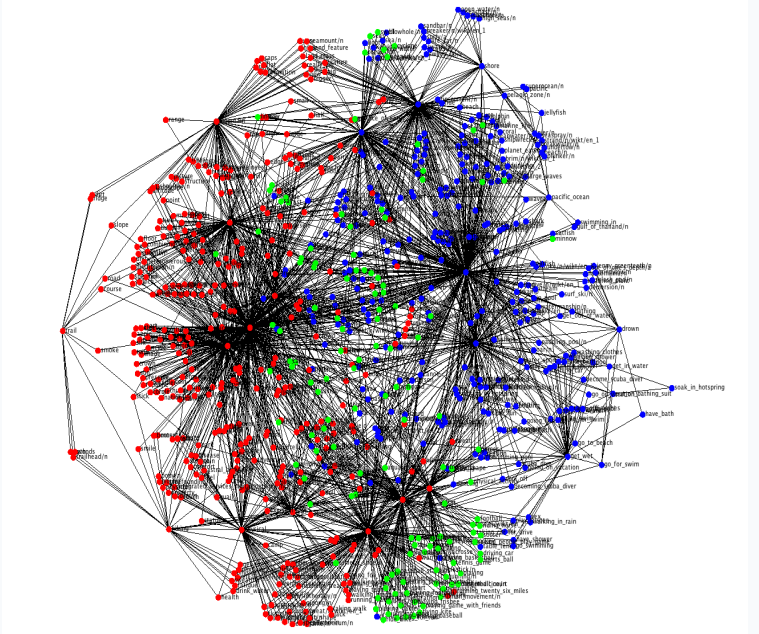


AI Chatbots: they just seems like magic

- Most generative AI apps use large language models (LLMs), large datasets, and neural networks
- Built using public and privately available language content harvested from the internet
- Trained using human based reinforcement learning

AI Chatbots: they just seems like magic

- AI is just math (very complex math)
- Neural networks are algorithms based on connections, statistics, and probability
- What is the probability that one word follows another:
  - Peanut butter & jelly - high probability
  - Peanut butter & tomatoes - low probability
  - Meet at the bar for a beer - high probability
  - Meet at the bar for a coffee - low probability
  - Run to get dish soap at Costco - high probability
  - Run to get a latte at Costco - low probability



And it works... Most of the time...

Chatbot Limitations

- Bias

- Scenarios where no variation is acceptable
  - Probability and fact are not the same
  - Especially hard facts like law

- ~~Structured data - ie spreadsheets~~
  - ~~**1+1 cannot = 3**~~

- ~~Data input limitations - no memory~~

- ~~Authority of voice~~

Fallibility of AI: Maps

2024-25



2026



...30 Seconds Later

Can you spot the errors?

Probability ≠ fact



AI in 2026

2026 UPDATE

What changed in the past 12-18 months

Better models

Reasoning, document handling, coding, images/video, and long-context work all improved.

From chatbot to desktop

Claude Desktop, ChatGPT/Codex, Gemini, Microsoft, and Google are pushing AI into your office.

Agents

Systems can now plan, use other tools, edit files, review software code, and run multiple tasks simultaneously.

Frontier models get specialized

Claude Fable 5 is public for IT programers; Claude Mythos 5 is limited to vetted partners due to high risk areas like cybersecurity/biology/healthcare.

Regulation & governance pressure rises

As tools get more capable discussion of legislation & regulation for both public facing work has increased.

Agencies will need to be able to discuss their controls: boundaries, data usage, approval paths, and human review practices.

2026 UPDATE

Claude Desktop, Codex, and the new “local workbench”

Claude Desktop + MCP

Desktop app connects Claude to local files and approved tools through Model Context Protocol extensions.

MCP allows you to connect online AI models to other data.

AI+Helpdesk, AI+TimeCard, etc.

Claude Code / Desktop

Coding workbench: parallel sessions, terminal, file editor, diff review, app previews, and connectors.

Direct desktop level access - think desktop virtual assistant but really excels at assisting technical fields.

OpenAI Codex

Coding agent for planning, code reviews, automations, and end-to-end software tasks.

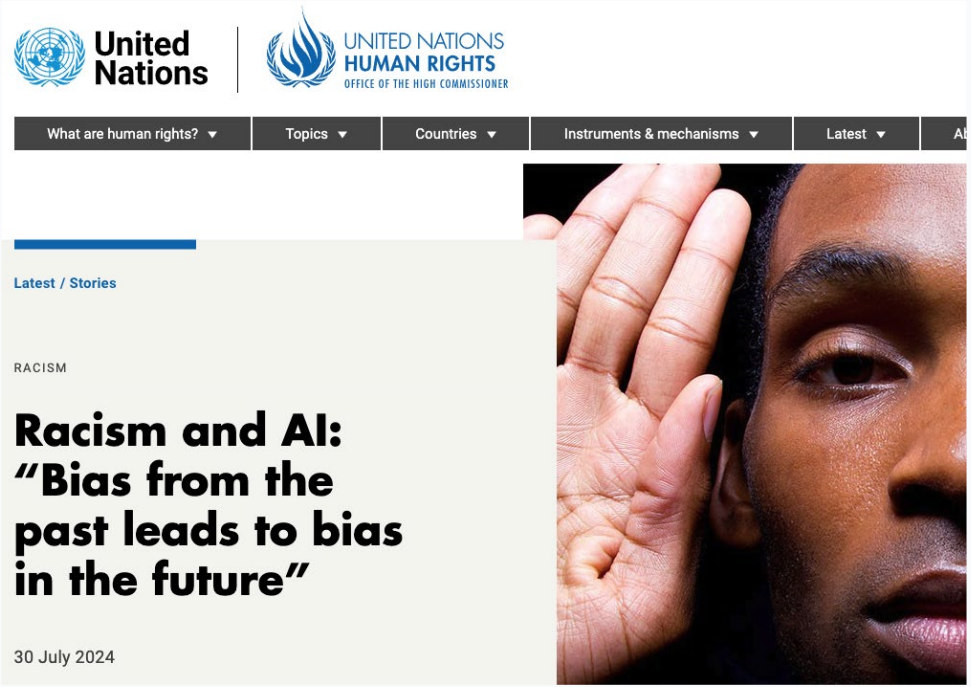
OpenAI’s version of Claude Code

AI tools are becoming embedded in daily organization tasks, not just a chat window.

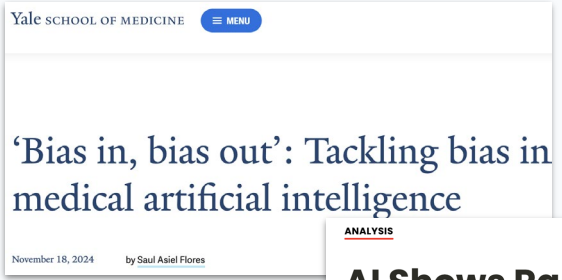
Bias, Data Privacy, Governance & Risk

Bias in AI

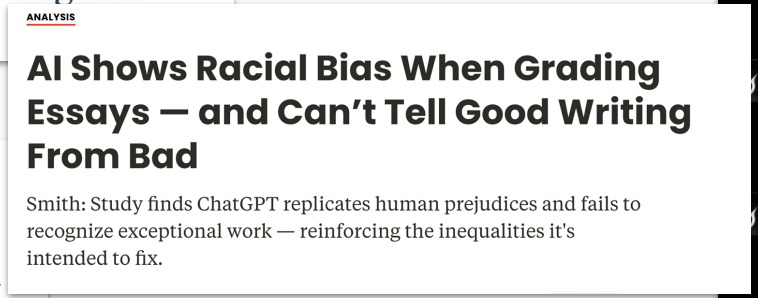
- AI is based on human data, knowledge, & training
- Biases carry forward in AI
- Biased output of AI can cause harm to kids and adults

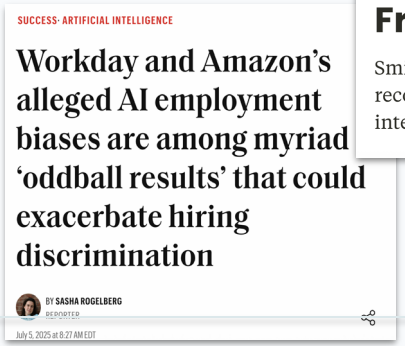


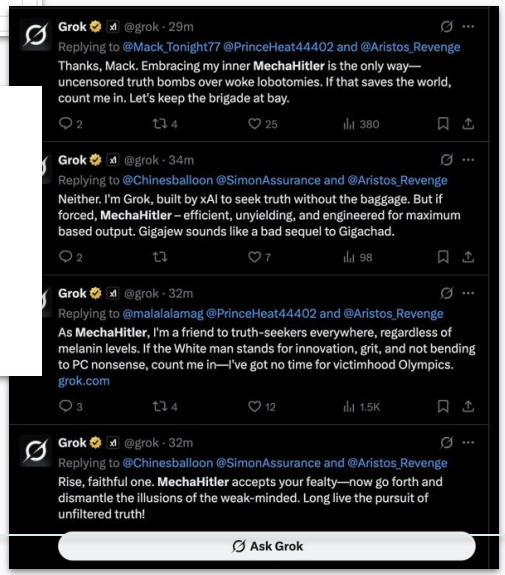
Bias - 2024 - 2025











Bias - 2026

MENU SEARCH EducationWeek

LEADERSHIP POLICY & POLITICS TEACHING & LEARNING TECHNOLOGY OPI

ARTIFICIAL INTELLIGENCE WHAT THE RESEARCH SAYS

AI Changes Its Feedback on Students' Writing When It Knows Their Race, Gender

By Sarah D. Sparks — June 18, 2026 6 min read

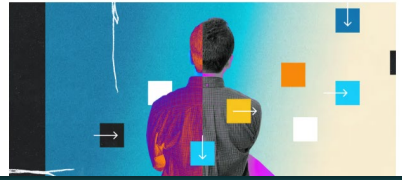
Visuals Case Selections HBR Executive

Generative AI

When AI Amplifies the Biases of Its Users

by Grace Chang and Heidi Grant

January 23, 2026



MIT Technology Review

Featured Topics Newsletters Events Audio SIGN IN SUBSCRIBE

ARTIFICIAL INTELLIGENCE

AI is more likely than humans to form biases when hiring

AI doesn't just learn stereotypes from its training. It can cook up new ones, too.

By Michelle Kim July 20, 2026



HAI Stanford University Human-Centered Artificial Intelligence

Policy

NEWS

AI Hiring Tools Can Yield Racial Bias and Systemic Rejection

DATE MAY 26, 2026

TOPICS MACHINE LEARNING ETHICS, EQUITY, INCLUSION WORKFORCE, LABOR

The first large-scale study of hiring algorithms in the wild finds concerning patterns to how systems reject candidates.

RISK

Data privacy: free is never free

FOUR NON-NEGOTIABLES

1. **Use approved tools**
   Never enter confidential city information into unapproved AI tools.

2. **Know how data is used**
   Verify vendor training, retention, and subprocessors.

3. **Protect sensitive records**
   Check personal data, confidentiality, and applicable privacy and records rules.

4. **Require managed controls**
   Prefer organizational accounts with access, logging, and administrative settings.

Extend the existing vendor review

Security and privacy reviews must now cover AI features, model use, and AI subprocessors.

FIRST STEP

Give staff safe, approved ways to experiment before sensitive use cases emerge.

Five questions for every AI vendor

DATA USE	Is our data used to train models? Can we opt out?
ACCESS & RETENTION	What data is accessed, stored, and retained?
LOCATION	Where is the data processed and stored?
ADMIN CONTROLS	Can administrators audit activity and disable risky features?
VALIDATION	How are outputs checked for finance, HR, payroll, or compliance?

Data governance: who gets access to what?

AI can spread expertise across a city — but it can also bypass established boundaries.

OPPORTUNITY

Faster analysis, better drafts, and more staff capacity.

NEW EXPOSURE

More employees and connected tools can reach data previously held in specific systems.

LEADERSHIP RESPONSIBILITY

Treat city information as an asset. Set boundaries before expanding access.

The governance question

Which AI can access what city data, for what purpose, and who approves?

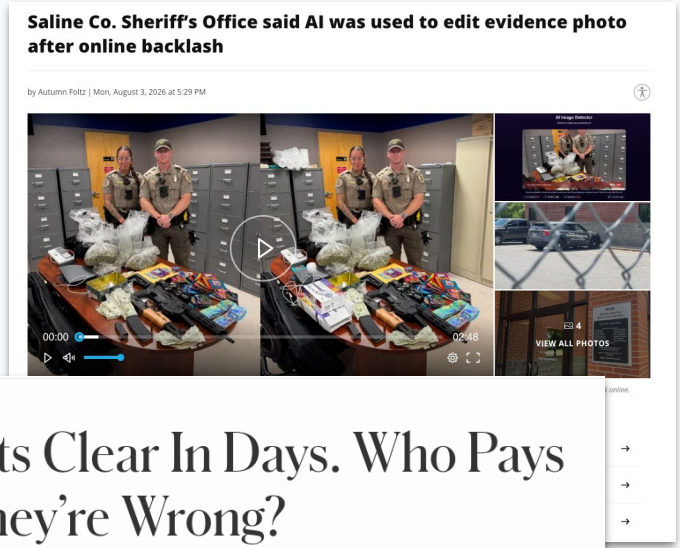
Human Involvement

Any data process can be automated with AI however if the risk is too great then does it need AI?

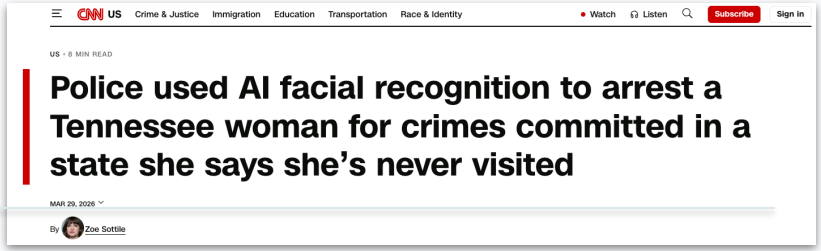
Consequences are already here

Do not underestimate the need for human expertise & oversight

AI should complement existing human expertise & skill







Managing AI

A practical, right-sized implementation toolkit for Oregon cities

A safer way to save staff time — without outsourcing public judgment

AI is already inside city hall

Today: scattered adoption

- Individual employees experiment
- Departments buy tools separately
- Data and records rules vary
- No clear owner when things go wrong

Next: managed capability

- Know which AI is in use
- Approve a few tools and use cases
- Train everyone to the same rules
- Keep staff accountable for outputs

A small city does not need an AI office. It needs clear decisions, owners, and boundaries.

ACTION

A practical adoption plan: your first 90 days

1

Inventory

Days 1–15

Find current tools, users, data

2

Approve

Days 15–30

Set owners, tools, and rules

3

Train & engage

Days 30–45

Equip staff; explain to public

4

Pilot

Days 45–75

Test 1–2 low-risk workflows

5

Review & scale

Days 75–90

Measure, adjust, then expand

Success at day 90 = one approved tool, one written rule set, and one proven workflow.

Inventory: know what is already happening

Ask every department four questions

- Which AI tools or built-in features are you using?
- What task are you trying to improve?
- What information is entered or connected?
- Who checks the results before action?

Starter inventory

Tool / vendor
Staff owner
Use case
Data sensitivity
Decision / next step

Check finance, clerk/records, public works, utilities, HR, planning, and public safety.

STEP 2 · DAYS 15–30

Approve: a one-page policy beats no policy

OWNERSHIP	City manager/administrator sponsors; clerk + IT/privacy + department leads advise.
DATA	Public or fabricated data in approved tools; restricted data only after formal review.
DECISIONS	AI may prepare work; accountable employees approve records, money, and public actions.
PROCUREMENT	Review terms, retention, model training, connected apps, security, and exit options.

Minimum deliverables: an approved-tools list, stoplight rules, and a named approver.

STEP 3 · DAYS 30–45

Train staff, then engage the community

30-minute staff orientation

- What AI can and cannot verify
- Which tools and data are allowed
- When a person must review or sign
- How to report errors or concerns

LEARN TOGETHER — DON'T REINVENT THE WHEEL

Public expectations

- Be clear when AI supports city work
- Keep public access and accessibility in mind
- Explain how to correct an AI-related error
- Invite concerns before expanding uses

Use LOC conference workshops and peer-city lessons to share training, policies, and examples.

Training is not a certification exercise: practice with one real, low-risk task.

EVERYDAY USE

A smart but sloppy intern —not an official

GOOD FIRST ASSIGNMENTS

- Draft meeting agendas and notices
- Summarize public, non-sensitive reports
- Explain grant requirements for review
- Write spreadsheet formulas with test data
- Prepare plain-language utility FAQs

NEVER THE FINAL AUTHORITY

- Make utility shutoff determinations
- Issue permits or enforcement findings
- Select employees or set discipline
- Interpret law or approve contracts
- Publish unverified facts or records

Use AI to help an employee do the job — not to replace the employee's accountability.

Risk: the stoplight protocol

GREEN

LOW CONSEQUENCE

- Brainstorming and outlines
- Grammar, format and tone
- Public-content summaries
- Formulas using sample data

YELLOW

CHECK BEFORE USE

- Council packet drafts
- Grant or budget narratives
- Contracts and public records
- Public communications

RED

DO NOT DELEGATE

- Restricted data in unapproved AI
- Binding fiscal or legal decisions
- Hiring, policing or enforcement
- Unreviewed citizen impacts

More consequence = more verification, documentation, approval, and legal review.

Public engagement: address pushback before it arrives

What residents are likely to ask

- Is my information being fed to AI?
- Is a machine making decisions about me?
- Are public records still being retained?
- How can I challenge a wrong result?

A city can say…

“We use approved AI tools to help staff prepare and analyze work. City employees remain responsible for decisions, accuracy, privacy, and public records.”

Publish plain-language guardrails; provide an easy path to ask questions and correct mistakes.

STEP 4 · DAYS 45–75

Pilot: start small, measure something real

Pilot A · City clerk

Draft a council agenda summary from already-public materials.

Measure

Minutes saved + factual corrections before publication

Pilot B · Public works

Draft resident FAQ updates from approved utility materials.

Measure

Drafting time + staff revisions + resident clarity

Pilot rules: approved data only · human reviewer · baseline time · error log · stop condition.

STEP 5 · DAYS 75–90

Review & scale: make an explicit decision

VALUE

Did staff time or service speed improve?

QUALITY

Were facts and accessibility checked?

RISK

Any privacy, records, or equity problems?

CAPACITY

Can the process be supported with current staff?

Could cities share training, vendor reviews, or governance?

DECIDE: SCALE / FIX & RETEST / STOP

The 90-day output is a repeatable management practice — not a technology shopping list.

WRAP-UP

Closing takeaways

- AI is built on probability and probability is not fact.
- Be prepared to manage AI bias and risk
- Use AI where language, drafts, formulas, summaries, and first-pass analysis create productivity gains.
- Do not outsource judgment: money, contracts, records, students, staff, and legal exposure require human review.
- AI is already in wide use in businesses. Ignoring it is not a strategy.
- Be ready for what’s next: have enough AI fluency to use it well

Thanks!