



AI's Rapidly Evolving Impact On Public Sector HR

California Public-Sector HR | 2025

November 20, 2025





Who We Are

►Colin Tanner

Colin J. Tanner is a founding partner of Aleshire & Wynder and Chair of the firm's Labor & Employment Practice Group. In these capacities, Mr. Tanner advises the firm's clients regarding their labor and employment practices as well as defends and/or prosecutes litigation claims on their behalf. He also practices in the areas of preventive liability, insurance coverage, and business litigation for the firm's public and private sector clients. Mr. Tanner has also authored or co-authored numerous articles regarding public sector labor and employment issues. Mr. Tanner received a B.A. in English & American Literature and Political Science from the University of California San Diego in 1987 and his J.D. from the University of California Hastings College of Law in 1990.

►Jude Poole

Jude Poole is the Director of Computing Services for the School of Biological Sciences at UC San Diego. In that capacity he has overseen an IT organization providing custom programming and database services in support of administrative services in the areas of Human Resources, Academic Personnel, Fiscal Services and Student Support. He chairs the Academic Information Technology Advisory Group and is the lead architect and programmer for the main campus Staff Performance Appraisal system and the Faculty FTE tracking and planning system. Mr. Poole received a B.A. in Political Science from UC San Diego in 1988.





Disclaimer

The information presented during this session is for informational and educational purposes only. The content is general in nature and may not reflect current legal developments or apply to your specific situation. Nothing presented here should be construed as legal advice or an offer of representation, nor is it a substitute for obtaining legal counsel on any subject matter.

An attorney-client relationship is not formed or intended to be formed by your attendance at this presentation, your review of these materials, or any communication related to this session. You should not act or refrain from acting on the basis of any information included in this presentation without seeking the appropriate legal or other professional advice on the particular facts and circumstances at issue from an attorney licensed in the relevant jurisdiction. The presenters and their affiliations disclaim all liability with respect to actions taken or not taken by a recipient based on any or all of the information contained in this presentation.






Billionaire investor Marc Andreessen says AI destroying jobs and making everyone poor is a ‘fallacy’—and even if that did happen, prices would drop





If any company or group, anywhere on the planet, builds an artificial superintelligence using anything remotely like current techniques, based on anything remotely like the present understanding of AI, then

**EVERYONE, EVERYWHERE
ON EARTH, WILL DIE**







Agenda



Quick review of 2024 foundations from last years presentation



What's new: models & tools



Agentic AI



Case studies & best practices



Legal & policy updates




Action plan & next steps





Quick Review: 2024 Foundations

- Job-related & validated: AI should assess legitimate skills and qualifications.
- Human oversight: Keep people in the loop to override or veto machine decisions.
- Monitor adverse impact: Regularly watch out for unintended bias and disparate impact.
- Data governance & privacy: Use only necessary, high-quality data; secure sensitive information.
- Reasonable accommodation: Provide alternatives for applicants and employees with disabilities.
- Procurement & recordkeeping: Follow public procurement rules and document decisions.





New Models & Tools (Nov 2024 – Sept 2025)

Rapid advances since late 2024 have delivered bigger context windows, specialised reasoning models and unified, multimodal architectures.

Key releases include:

Model	Release	Highlights
GPT-5	Aug 7 2025	Unified reasoning+speed; agentic tasks; real-time router
O3 series	Jan–Jun 2025	Reflective reasoning; high/medium/low effort levels; precision & math
Claude 4.5	Sept 29 2025	State-of-the-art coding; 30h autonomy; lower prompt-injection risk
Llama 4 (Scout/Maverick)	Apr 5 2025	Multimodal; 10M-token context; outperforms GPT-4o/Gemini
Mixtral 8×22B	Apr 17 2024	Mixture-of-experts; 64k context; open-weight & efficient
DeepSeek V3/R1/V3.2-Exp	Dec 2024 – Sept 2025	MOE models; RL-enhanced reasoning; Sparse Attention
Gemini 2.5	Mar 2025	Native multimodality; 1M-token context (2M soon)
OpenAI Atlas browser	Oct 2025	Chromium with built in ChatGPT





The Scale of Modern AI

Company	Approx Market Cap/Val (oct 2025)	Business
Nvidia	\$4.6T (market cap)	AI compute: GPUs, systems, networking, platform
OpenAI	\$500B (valuation)	Frontier models & apps (ChatGPT)
xAI	~\$200B (reported; disputed)	Frontier models (Grok); X data; infra buildout
Anthropic	\$183B (valuation)	Frontier models/safety (Claude)
Meta	\$1.8T (Market Cap)	Social Network, AI Models (Llama)
Alphabet	\$3T (Market Cap)	Google, etc, AI Models (eg Gemini)
Microsoft	\$3.8T (Market Cap)	Enterprise and consumer software, AI Cloud





CORE PRACTICAL SKILL: Prompting for HR tasks

- State role, task, audience, jurisdiction
- Supply policy text or excerpts to ground answers
- Define output format & length upfront
- Include 1–2 examples and a scoring rubric
- Require citations & “say I don’t know” when unsure
- Protect privacy: no PII (Personally Identifiable Information) —use mock data





Workflow & risk controls in prompts

- Keep human review—no auto-decisions
- Prefer retrieval over pasting; whitelist sources
- Add bias checks & inclusive-language passes
- Version prompt templates; record model & settings
- Records-ready: save prompt, output, sources
- Red-team with edge cases & prompt-injection tests





Drop in Prompt Template (cut/paste)

- **ROLE:** You are a neutral HR analyst for a [jurisdiction] government agency.
- **TASK:** [What to produce] for [audience/purpose].
- **CONTEXT:** Use only the provided excerpts from [named, approved sources].
- **CONSTRAINTS:** Follow [policy names/sections]. If info is missing, write "Unknown based on provided text."
- **OUTPUT:** [exact format: bullets/table/headings], ≤[length].
- **QUALITY BAR:** Match this example's tone and structure: [short example]. Apply rubric: accuracy 50, clarity 30, completeness 20.
- **CITATIONS:** Bracket the exact section (e.g., "[SX.Y]") after each claim.
- **SAFETY:** Do not include or infer PII. Use synthetic placeholders.
- **REVIEW:** List 3 items the HR reviewer must verify before use.





Defining Agentic AI



Reasoning & Decision-Making

Plans multi-step tasks, generates hypotheses and evaluates options. “Thinking” models are an example



Tool Use

Invokes APIs, databases, search and code to act on the world



Memory & Context

Stores context and history to inform actions and avoid repeating mistakes

Why it matters for HR:

Agentic AI enables highly autonomous workflows—like automatically sourcing, interviewing and selecting candidates. While powerful, these systems can amplify bias or leak sensitive data if not carefully governed. Maintain audit trails, limit tool access, and ensure every autonomous flow has human checkpoints. All users should be VERY careful about granting agentic AI direct access to local systems.





Case Study: Resume Screening & Ranking

Scenario:


A public agency uses a model to screen resumes for a Customer Service role. Applicants are ranked according to scores derived from skills, experience and education.

Benefits:

Accelerates processing; consistent scoring across hundreds of resumes; focuses human reviewers on top candidates.

Risks & Mitigations:

- Bias in training data may rank certain groups lower.
- Proxy variables (e.g. university name, addresses) can reflect socioeconomic status.
- Always have a human review and override rankings.
- Run regular adverse impact tests and adjust scoring weights.





Mobley vs. Workday Inc.

- What Mobley alleges (and why HR should care)
- Procedural posture & key rulings (7/12/2024 MTD order)
- Direct-liability “agent” theory allowed; “employment agency” theory rejected
- Claims surviving vs. dismissed (disparate impact survives; intentional discrimination/FEHA aiding-abetting narrowed)
- 2025 update: preliminary ADEA collective certified (40+ applicants; “denied employment recommendations”)
- Why this matters to public employers using ATS/AI
- California’s new AI/ADS rules (effective Oct 1, 2025): duties & recordkeeping
- Practical to-dos: inventory, contracts, audits, documentation, accommodation, human oversight





Case Study: Interview Question Generation

Scenario:


An AI assistant drafts tailored interview questions based on job descriptions and applicant CVs.

Considerations:

- Questions must be job-related and uniformly applied.
- Avoid prompts that probe protected traits (age, race, disability).
- Ensure the assistant doesn't reveal or infer confidential information.
- Provide candidates with accommodations (e.g. different formats).

Best Practices:

- Pre-approve question templates with legal counsel.
- Combine generative questions with human insight.
- Offer a mechanism for human follow-up and clarification.





Case Study: Performance Appraisal & Coaching

Scenario:

An AI system analyses productivity metrics, feedback and attendance records to provide performance ratings and coaching recommendations.

Considerations & Mitigations:

- Transparently communicate how data is collected and used.
- Ensure metrics are job-related and not proxies for protected characteristics.
- Provide employees with the opportunity to contest scores and add context.
- Use AI insights as one data point; final decisions rest with managers.
- Retain performance data for four years as required by FEHA.





Emerging Legal Landscape: FEHA AI Regulations

Effective Oct 1 2025:

California's Civil Rights Council clarified that using an automated decision system (ADS) may violate FEHA if it harms applicants or employees based on protected characteristics. Employers must:

- Maintain ADS-related records (datasets, scoring outputs, audit logs) for at least four years;
- Treat games, quizzes or puzzles that elicit disability information as medical inquiries;
- Ensure any automated decisions are free from discrimination based on race, gender, disability or other protected traits;
- Define and understand “automated decision system”, “agent” and “proxy” in vendor contracts.

See articles such as: Don't Let AI Put Your Agency at Risk—FEHA Compliance in Hiring Explained

<https://www.calpublicagencylaboremploymentblog.com/employment/dont-let-ai-put-your-agency-at-risk-feha-compliance-in-hiring-explained/>





Compliance Checklist (FEHA)

- ✓ Audit AI tools used in hiring, promotion, evaluation and discipline.
- ✓ Request anti-bias protocols and data-use statements from vendors.
- ✓ Conduct bias testing with legal guidance pre- and post-deployment.
- ✓ Update recordkeeping policies to preserve ADS data for four years.
- ✓ Ensure every AI-assisted decision includes human oversight.
- ✓ Train HR and managers on new definitions and obligations.





Emerging Legislation: SB 53 (Effective Jan 1, 2026)

Transparency in Frontier Artificial Intelligence Act

Applies to developers of large “frontier” models (>\$0.5B revenue). Requires:

- Publish safety & security protocols on their websites;
- Disclose testing procedures and mitigation of catastrophic risks;
- File transparency reports before releasing new or updated models;
- Allow whistleblowers to report critical safety incidents without retaliation;
- Coordinate with the state to develop CalCompute – a public cloud cluster for safe AI research.

Why it matters for HR: Procurement & vendor selection will depend on whether vendors comply with SB 53. Agencies must request transparency reports and monitor safety disclosures.





Emerging Legislation: SB 7 [VETOED]

No Robo Bosses Act (VETOED)

Vetoed October 2025. If enacted, it would have:

- Prohibited employers from relying solely or primarily on automated systems to discipline or terminate workers;
- Required disclosure to employees when AI influences high-impact decisions;
- Provided employees with access to the logic of the system and the ability to contest decisions;
- Mandated documentation of the reasons for any AI-assisted disciplinary action.

Implication: Even without SB 7, employers face growing risk under FEHA AI regulations, ADEA/Title VII litigation (eg Mobley) and AG enforcement. Treat SB 7 as a preview of concepts likely to reappear in future legislation, CBAs, or local policies.





Federal Signals: OMB AI Policies (Apr 2025)

OMB M-25-21 (Use of AI):

- Encourages accelerated AI adoption by removing bureaucratic barriers.
- Mandates agencies designate a Chief AI Officer and create cross-functional AI Governance Boards.
- Requires pre-deployment testing, AI Impact Assessments and ongoing human oversight for high-impact use cases.
- Calls for annual public inventories of AI use cases and waivers.

OMB M-25-22 (AI Procurement):

- Promotes acquisition of U.S.-developed AI solutions and talent.
- Requires procurement contracts to restrict vendor use of government data for training.
- Emphasises cross-functional engagement and risk management in AI procurement.





Action Plan: First 90 Days

- Inventory existing AI & ADS use in your HR processes.
- Review vendor contracts for transparency & bias testing obligations.
- Launch bias audits on high-impact ADS (screening, appraisal).
- Develop an internal AI governance policy & designate a responsible officer.
- Provide training on FEHA regulations and federal AI policies.
- Engage IT & legal teams to monitor agentic AI adoption and tool selection.





Case Study: Succession Planning & Development

Scenario:

A predictive model identifies employees with leadership potential and suggests personalised development plans.

Opportunities:

- Supports equitable talent mobility and reduces turnover.
- Tailors training resources to individual strengths.

Risks & Safeguards:

- Models may unintentionally favour employees who already fit leadership stereotypes.
- Provide transparency into recommendation logic and allow employees to opt in.
- Combine AI insights with manager feedback and 360-degree reviews.
- Regularly update training data to reflect evolving competencies.





Looking Ahead: Trends & Challenges

- Multi-agent systems will automate end-to-end HR tasks (from sourcing to onboarding).
- Longer context windows enable comprehensive talent analytics and personalized HR support.
- Native multimodal models (text/audio/video) will analyse interviews, training & communications.
- Increased regulation will demand greater transparency, auditing and data retention.
- AI literacy becomes an essential skill for all HR practitioners.





Conclusion & Questions


AI can transform HR processes, but success hinges on fairness, transparency and accountability.

Key takeaways:

- Invest in responsible AI governance and bias audits.
- Stay abreast of emerging laws and actively engage legal counsel.
- Embrace agentic tools thoughtfully: empower staff, avoid over-automation.

We look forward to your questions and dialogue!







Notes, References, Resources

Books

- ❖ Co-Intelligence: Living and Working with AI – Ethan Mollick
- ❖ Weapons of Math Destruction – Cathy O’Neil
- ❖ If Anyone Builds It, Everyone Dies – Eliezer Yudkowsky and Nate Soares
- ❖ What is ChatGPT Doing... And Why Does it Work – Stephen Wolfram
- ❖ AI Snake Oil – Arvind Narayanan and Sayash Kapoor





Notes, References, Resources

Misc

- ❖ Grok System Prompt: https://github.com/xai-org/grok-prompts/blob/main/grok4_system_turn_prompt_v8.j2
- ❖ ChatGPT 800 Number: 1-800-CHATGPT (30 mins/month)
<https://help.openai.com/en/articles/10193193-1-800-chatgpt-calling-and-messaging-chatgpt-with-your-phone>
- ❖ How difficult is it for a programmer (or anyone) to run AI locally? Not Very: <https://dzone.com/articles/zero-to-local-ai-ollama-python>
- ❖ AI Powered Fake Job Applicants (and how to spot them)
<https://ogletree.com/insights-resources/blog-posts/cut-the-ai-bait-and-switch-tips-for-employers-to-spot-fake-job-applicants/>

