California Community Colleges

MEMORANDUM

November 10, 2026

ISTI 25-300-03 | Via Email

TO: Chief Executive Officers, General Counsels, Chief Instructional Officers, Chief Student Services Officers, Academic Senate Presidents, Chief Technology Officers, Chief Business Officers, and Chief Human Resources Officers

FROM: Craig Hayward
Vice Chancellor, AI Strategy and Digital Transformation

RE: AI and Privacy
AI Guidance Series

Introduction

The latest installment in the ongoing AI guidance series, this memorandum builds on the guidance in ISTI 25-300-01, ISTI 25-300-02 and ESS 26-52 to support the advocacy for and protection of all California Community College constituents' privacy. Grounded in the **HUMANS** Framework and the Vision 2030 commitment to equity, this document offers recommendations on how institutions can protect constituents' privacy in the age of artificial intelligence (AI).

By engaging with this memo's guidance collectively, the California Community Colleges can better uphold our values of privacy, agency, and success for students, faculty, and administrators.

How to Use this Memo

Protecting and advocating for the privacy of students, faculty, staff and administrators is integral to the California Community Colleges' mission of equity, access, and student success. While AI technology brings great promise of improving student outcomes and transforming teaching, learning, and student services, its capability to dramatically increase the scale of sensitive data collection and potential for exploitation requires institutions to revisit how existing protections extend into the AI-enabled world and where new measures are necessary.

Every college's administrators, faculty, staff, and students are wrestling with questions surrounding privacy and AI. This memo is written primarily for administrators and institutional leaders as a resource to supplement these ongoing campus conversations.

This memo meets institutions at a variety of maturity levels in addressing privacy concerns related to the use of AI. As AI functionality and tools are constantly changing, an essential part of this work is continually reviewing the privacy efforts each institution already has underway. We look forward to

your feedback and continued partnership as we collaboratively define how and when to integrate AI, while ensuring its responsible and informed use.

This memo outlines important privacy risks and protections, critical institutional actions, and privacy rights for students and employees. The following companion resources exist on the AI microsite:

1. **Institutional Privacy Risk Assessment** intended to inform institutional leaders' agendas with regard to privacy. These recommendations are not comprehensive but can serve as a guide to assess your college's progress in addressing AI privacy issues.
2. **Professional Responsibilities** of employees of the California Community Colleges to protect the privacy of other constituents and recommendations to operationalize these practices. This section is similarly intended to provide administrators and institutional leaders with substance to reuse in their local contexts.
3. **Individual Privacy Rights** of California Community College constituents and steps individuals can take to exercise these rights. This section is intended to be a template for administrators and institutional leaders to adapt and deploy locally to help address questions and concerns of students and employees.

This memo does not constitute, nor is it intended to substitute for, legal advice. Districts and colleges should consult their own legal counsel with specific questions about privacy and any legal implications for their policies and practices.

Background: The Privacy Landscape

The right to privacy encompasses many related freedoms: freedom from intrusion, control over information, and a right to personal autonomy. The California Constitution includes privacy as an inalienable right for all Californians.[1] Historically, laws and policies protected these rights by controlling what information gets collected, by whom, and for what purpose. The mechanism of these protections leaves gaps in the AI era; even if sensitive information is not collected explicitly, AI tools can infer sensitive information from other data.

In the higher education context, protecting privacy means ensuring individuals retain control and agency over their personal information and records, proprietary works, identity, and decision-making processes. Practically, institutions must govern what information AI tools can collect and retain and

[1] Cal. Const. art. I, § 1.

how AI inferences can be used and shared. It is critical to be transparent about when personally identifiable information (PII) can legally be shared with and without consent. Doing so also depends on foundational data governance: knowing what data the institution holds, how it is classified, who owns it, and where it flows.

This memo summarizes some current best practices and provides recommendations for protecting privacy in the age of AI. A landscape of overlapping state and federal laws, industry standards, and district and institutional policies and practices protects individual privacy. Governing laws and policies may also depend on who is handling an individual's information, what that information is, and for what purpose it was collected and is being retained. While some of these laws and standards have been revisited to address new AI technologies, many still lack specific application guidance regarding AI. For example, a college using AI to make a hiring decision, a faculty member using AI when grading a student, and the third-party AI tool itself are each governed by laws providing different protections.

Examples of each category include:

- **Federal Laws:** At the time of writing, there are no federal privacy laws directly and explicitly pertaining to AI, but laws such as Family Educational Rights and Privacy Act (FERPA) and Health Insurance Portability and Accountability Act (HIPAA)[2] do apply to an AI context.[3] Additionally, the Gramm-Leach-Bliley Act (GLBA) mandates that institutions that provide student aid maintain a formal Information Security Program establishing administrative, technical, and physical safeguards for how data is collected, stored, and protected. Because GLBA compliance is evaluated during annual institutional audits any AI tool integration must align with these required security controls.
- **California Laws:** California does have laws explicitly governing AI as well as additional laws that apply in an AI context. These include, but are not limited to, the California Consumer

[2] While HIPAA does provide privacy protections for health information in particular cases, it does not broadly apply to the California Community Colleges and is therefore not covered in detail in this memo. See the Joint Guidance on the Application of FERPA And HIPAA to Student Health Records for more information on how HIPAA applies.

[3] While there are Executive Orders (EOs) that pertain to AI, as a rule federal EOs only direct the actions of Federal agencies, not higher education institutions.

Privacy Act (CCPA),[4] California Invasion of Privacy Act (CIPA), California Ed. Code §87031, and the recently enacted Higher Education Student Information Protection Act (HESIPA).[5]

- **Institutional Policies:** Institutional policies at the California Community Colleges mandate vetting measures for technology services and tools (including AI), providing another layer of protection. For example, vendor confidentiality and data protection clauses are included in the CollegeBuys program contracts to define restrictions on vendor use of student and employee data. District-level privacy policies and practices further protect student and employee data and the **HUMANS** principles, particularly **M -** "Managed Privacy Controls" and **N -** "Notice and Explanation," promote development and use of privacy-preserving AI technologies.[6]
- **Industry Standards:** The International Organization for Standardization (ISO) and the National Institute of Standards and Technology (NIST), among other organizations, have published frameworks on privacy standards and best practices. Notable frameworks include: ISO/IEC 27701 and 42001, NIST's AI Risk Management Framework, NIST Privacy Framework.

Privacy laws, policies and standards have existed and protected community college constituents for decades. While some recent measures provide privacy protections tailored to an AI-enabled world, much of the privacy landscape does not. This memo operationalizes the latest thinking on privacy in the age of AI and provides guidance on applying principles of existing law to today's technology.[7]

- Vendor vetting details will be further explored in Part 4 of the AI Memo Series.
- Privacy rights as they relate to students younger than 18 and dual-enrolled in K12 and higher education will also be outlined in separate guidance, as these students (and responsibilities for their parents and guardians) are impacted differently by privacy laws than the laws that govern the privacy rights of those over 18.

Privacy Risks and Protections

It is critical for colleges to clearly communicate and provide professional development opportunities to students, faculty, staff and administrators on potential risks about what information, including

[4] While CCPA does not apply directly to districts, it does apply to for-profit businesses and vendors that districts contract with.

[5] The Higher Education Student Information Protection Act (HESIPA) is set to become operative on July 1, 2027; several bill proposals are also currently working their way through the California legislature.

[6] **HUMANS** Guiding Principles and Framework

[7] Note: these subjects will be further explored in greater detail in future memos.

personal information, may be collected through AI tools and why.[8] Traditional privacy law was built around discrete acts: a record disclosed, a file accessed, a consent form signed. AI does not fit neatly into these privacy frameworks: it can collect data continuously, at scale, across every interaction, and it can make inferences about individuals that go far beyond what they explicitly shared. It can use personal information to complete tasks and train future systems, often without users' knowledge.[9] AI is capable of performing consequential actions, which is why the **HUMANS** principles specify that humans must be in the loop. Below are examples of harms that could result from a privacy breach.

- **Personally Identifiable Information (PII), Sensitive Data, and Intellectual Property (IP) Exposure or Misuse** - AI tools often collect, store, and use data entered by individuals to train future models or share with third parties. When students, faculty, staff and administrators use AI, student records, employee information, and institutional data can be exposed or misused. Proprietary research, curricula, and other institutional IP shared with AI tools may also lose legal protections or be reproduced without authorization.
- **Compromised Identity or Consent, Deepfakes** - AI can be used to impersonate or create realistic "deepfake" images or videos of individuals without permission from just a single image or short snippet of video. As AI use grows, the risks of deepfakes, non-consensual intimate imagery (NCII), and AI impersonation increase. Scammers can use deepfakes of individuals to pass identity validation tests, raising risks of identity fraud.[10]
- **Behavioral Exploitation and Surveillance** - AI tools continuously collect vast amounts of data about how users interact, learn and behave. Unlike previous technologies, AI can use this data at scale to appear as a trusted friend, mental health confidant, or academic advisor in ways that exploit psychological vulnerabilities and foster unhealthy dependency. The same data can also be used to monitor and at times incorrectly flag student work products or behavior, creating a chilling effect on AI engagement and literacy.

Existing privacy frameworks, many of which were designed decades before generative and agentic AI existed, do not explicitly cover these harms. Individuals' data practices directly shape the threat

---

[8] "Personal information" can be defined in a variety of ways depending on the context and use case but generally includes fields such as full names, email addresses, phone numbers, social security numbers…etc. Clarify with your institution what data is included in definitions of "personal information".

[9] In addition to restricting collection and use of immigration status, sexual orientation, gender identity, and reproductive and sexual health information and adding deletion rights and a private right of action for higher education students, HESIPA also makes AI training on student data illegal beginning July 1, 2027; non-student employees are not protected by federal or state law in this regard.

[10] To be clear, distributing deceptive election-related deepfakes without clear disclosure is prohibited near elections (AB 730). Additionally, creating or sharing non-consensual sexually explicit deepfakes incurs civil liability and criminal penalties with heightened severity for cases involving minors (AB 621; SB 981).

landscape that campus technology, security, and privacy teams must defend. Using unapproved AI tools and vendors (aka 'shadow AI') and unnecessary data collection create visibility gaps and vulnerabilities that may compromise institutional security. Using institutionally approved AI tools allows colleges to benefit from enforceable vendor commitments, rigorous privacy safeguards, and ongoing oversight that unapproved alternatives cannot provide. A clear privacy imperative for institutions is providing the resources and training needed to build out and implement these capacities effectively.

Institutional AI Procurement

Students, faculty, staff and administrators generally access AI tools in three main ways:

1. Institutions procure tools directly, requiring colleges and districts to enter into institution-specific agreements with AI vendors that determine how user data is handled and how privacy is protected.
2. State or system-level procurements are made available to colleges and districts.
3. AI tools that were purchased or adopted for personal use are used for work without institutional review for data privacy and security (i.e., shadow AI).

As discussed above, shadow AI use is the riskiest means of AI access from a privacy perspective due to a lack of default privacy protections. Institutional procurement practices provide an opportunity to implement rigorous privacy protections. When procuring tools from AI vendors directly, institutions should review vendor privacy policies, terms of service and contract terms and conditions to understand how features work, with a particular focus on privacy-enhancing features to protect students (including minors), faculty, staff and administrator privacy. AI features are also increasingly added to already-licensed platforms through vendor updates, often enabled by default; these updates should also be reviewed even when no new procurement occurs.

- **Include privacy and information security into the contract.**[11] Because these terms and conditions are open to negotiation, institutions can secure protections that individual users never see. Chief among these are bans on using institutional data to train the vendor's models, firm retention and deletion schedules, tight limits on third party access, clear IP ownership, and liability assignment to vendors and FERPA-aligned terms wherever education records are in play.[12] Recommended data provisions include imposing requirements on data classification, US data residency, and data breach notification timelines. If agentic AI

[11] Beginning July 2027 certain student data privacy terms will be required of AI vendors as described in HESIPA.

[12] The next memo in this series will cover AI vendor vetting and contracts in greater detail.

capabilities are included, it is critical that strong policy and guardrails are also implemented to ensure that autonomous AI agents do not inherit broad user permissions through Single Sign-On (SSO) or otherwise that could inadvertently grant them rights to access and manipulate sensitive data such as protected educational records, financial aid systems, or employee HR data. Institutions must enforce strict permission scoping and active monitoring to prevent autonomous agents from driving unauthorized data exposure.

- **Govern and disclose how data is used internally.** Administrators managing institutional AI tools should set a high bar for who may access end-user data and why, draw a clear line between security monitoring and privacy intrusion, and tell and remind users plainly how their data is used, retained and deleted, where opt-outs are available, and how these terms differ from consumer versions.
- **Plan for public records, retention and legal holds.** Prompts, chat logs, transcripts and AI-generated outputs created in the course of district business may be public records, may be subject to records-retention schedules and litigation holds, and may be discoverable. Contract deletion schedules and individual deletion requests should be reconciled with these obligations in advance, and users should be informed that their interactions with AI tools on institutional accounts may not be private from the institution.

When considering opting in to **state- or system-procured AI tools,** local institutions should review vendor agreements prior to providing tool access to students, faculty, staff or administrators.

- **Ensure alignment with local policy.** Due to varying Board Policies, collective bargaining agreements and governance standards across districts, institutions should periodically review terms and conditions and consult legal counsel to ensure that AI vendor agreements do not violate local policy.
- **Clearly communicate privacy protections.** Institutions opting in to AI tools contracted at the system level should review and communicate the terms and conditions to users. Relevant details may include how user data is processed, used and retained; how intellectual property ownership is allocated; who is legally permitted to use the tool (e.g., minors cannot consent to the terms of service of certain AI tools); and any additional protections that the institution has in place to protect users, if applicable.

Institutional Actions

In order to adapt existing privacy frameworks to address and protect privacy considerations in the age of AI, institutions should consider focusing on the following three next steps. A more comprehensive list of best practices can be found in the **Institutional Privacy Risk Assessment.** These recommendations do not constitute, or substitute for, legal advice.

01. Ensure AI Governance Policies, Practices and Processes Are in Place	
Do ☑	Don't ✖
Integrate AI governance responsibilities across functions and into existing participatory governance with clear owners	Treat a standalone AI policy document as sufficient governance without processes and owners behind it across functions
Review committee membership to ensure appropriate representation across privacy, cyber, HR, legal, unions and senates	Assume existing IT or cybersecurity committees cover AI privacy by default
Review privacy risks and controls of all AI procurements before and regularly after deployment	Rely on ad hoc or informal review when individual departments adopt AI tools
Build AI governance atop data governance, with named data owners and stewards	Create AI governance structures without aligning with data governance structures

02. Perform a Privacy Risk Assessment to Prioritize Institutional Next Steps	
Do ☑	Don't ❌
Leverage the institutional risk assessment to identify strengths and weaknesses in existing institutional privacy program	Treat a one-time assessment as sufficient
Enhance AI vendor vetting criteria to include comprehensive privacy protections	Accept vendor privacy assurances or certifications in place of your own vetting
Use the findings of the risk assessment to prioritize faculty, staff and administrator capacity, budget allocation and timelines	Ignore outcomes of risk assessments or fail to implement appropriate changes
Establish a data classification scheme and inventory of AI tools and agents in use before negotiating vendor terms	Negotiate classification or retention terms the institution hasn’t defined for itself

03. Clearly Communicate AI Privacy Policies and Rights to Constituents	
Do ☑	Don't ❌
Communicate student, faculty, staff and administrator rights and responsibilities on AI and privacy clearly and accessibly	Bury AI privacy policies and practices in dense legal language or make them hard to find
Communicate faculty, staff and administrator professional development opportunities widely and accessibly	Assume one announcement reaches all students, faculty, staff and administrators
Include ongoing privacy education among faculty, staff and administrator training requirements	Omit privacy from faculty, staff and administrator training curricula or limit training to annual compliance sessions

Conclusion

Protecting an individual's privacy is an ongoing practice as the technology and legal landscapes continue to evolve. In addition to protecting privacy rights, it is critical that we continue to expand our understanding of AI technologies, the evolving risks associated with them, and best practices for developing and deploying them safely and securely. Doing so will help ensure innovation at the California Community Colleges remains both responsible, trusted and secure. In our next memo, we will explore how integrating privacy, security, and utility ensures a comprehensive AI vendor vetting process, and detail our engagement with the new intersegmental body on GenAI vendor and training standards established by AB 2392 (Fong). We look forward to hearing your feedback on this guidance and ways we can work together to support the needs of our students and employees in this era of AI-enabled higher education.

Additional Guidance and Resources

- AI Memo Series
    - AI Teaching and Learning Memo
    - AI Safety and Security Memo
    - AI Literacy Memo
- Frequently asked questions (FAQs)
- AI microsite
- Reach out to aiguidance@CCCCO.edu with any additional questions or concerns

Craig Hayward
Vice Chancellor, AI Strategy and Digital Transformation

cc:

Sonya Christian, Chancellor
Rowena Tomaneng, Deputy Chancellor
Chris Ferguson, Executive Vice Chancellor
John Hetts, Executive Vice Chancellor
James Todd, Vice Chancellor
Renee Craig-Marius, Vice Chancellor
Siria Martinez, Vice Chancellor
Jerry Deschler, General Counsel
Jory Hadsell, Executive in Residence
Don Daves Rougeaux, Senior Advisor
AI Council