



Palo Alto University (PAU) Artificial Intelligence (AI) Policy

Division Responsible: Finance and Operations

Office Responsible: Information Technology

Publish for Compliance:

Publication Date:

Last Review Date:

Original Creation Date:

Next Review Date:

Introduction and Purpose	2
Scope & Applicability.....	2
Roles and Responsibilities.....	4
Responsible Use of AI in Research & Administration.....	4
Reporting & Enforcement	8
Ethical Principles for AI Use	10
AI Subscription and Licensing.....	10
Policy Review.....	12
Appendix A: Definitions & Key Terms.....	13
Appendix B: Resources & References	14
FAQ:	15

Palo Alto University

Introduction and Purpose

Artificial Intelligence (AI), including Generative AI, is rapidly transforming how Palo Alto University (PAU) supports teaching, research, and administrative operations. These technologies present significant opportunities to enhance student learning, improve operational efficiency, and advance the university's mission of education, research, and community service.

At the same time, AI introduces important considerations related to data privacy, academic integrity, security, and regulatory compliance. The use of AI within the university must be governed in a manner that protects institutional data, ensures responsible use, and maintains trust across the PAU community.

This policy establishes clear expectations for the appropriate use of AI at PAU. It defines what is permitted, restricted, and prohibited when using AI technologies in academic, research, and administrative activities.

The objectives of this policy are to:

- Enable the responsible and effective use of AI to support academic and operational excellence
- Protect sensitive institutional data and ensure compliance with applicable regulations (FERPA, HIPAA, GLBA)
- Maintain academic integrity and uphold ethical standards in teaching and research
- Provide clear guidance to students, faculty, and staff on acceptable AI use
- Ensure appropriate human oversight, accountability, and transparency in all AI-assisted activities
- Protect clinical and health-related data — including therapy session transcripts, notes, and telehealth records — and support sound decision-making in activities that impact clients served through PAU-operated clinics

PAU is committed to leveraging AI in a way that supports innovation while safeguarding the university's core values of integrity, accountability, and respect.

Scope & Applicability

This policy applies to all members of the Palo Alto University (PAU) community, including students, faculty, staff, administrators, researchers, contractors, consultants, and affiliated partners who access or use PAU data, systems, or resources.

In Scope

This policy applies to the use of:

- Public AI tools (e.g., ChatGPT, Claude, Gemini, DALL-E, and similar platforms)
- AI capabilities embedded within enterprise systems (e.g., Microsoft 365 Copilot, Google Workspace, Jenzabar, Salesforce)

Printed copies of this document are uncontrolled. In the case of a conflict between printed and electronic versions of this document, the controlled version published online prevails.

Commented [DS1]: Not sure where to add this but in clinical work it is more than transparency, it is written consent

Commented [GS2]: Should we add this? Canvas LMS and third-party LTI integrations such as Turnitin and Respondus Monitor.

Palo Alto University

- Third-party or cloud-based platforms that process, analyze, or generate content using PAU data
- AI used on both university-owned and personally owned devices when conducting university-related work
- Learning management systems and integrated educational technologies (e.g., Canvas LMS and third-party LTI integrations such as Turnitin and Respondus Monitor)

This policy applies regardless of whether AI tools are:

- Free or paid
- Individually purchased or institutionally licensed
- Accessed directly or embedded within other software

Environments Covered

This policy applies across all environments where PAU systems or data are accessed, including:

- On-campus locations
- Remote or hybrid work environments
- Third-party or vendor-managed systems

Data and Compliance Requirements

All use of AI must comply with:

- Applicable laws and regulations (including FERPA, HIPAA, GLBA, and copyright law)
- Applicable accessibility standards, including Section 508 / ADA requirements and WCAG 2.1 AA guidelines, for AI tools integrated into teaching and course materials
- PAU policies, including but not limited to:
 - Academic Honor Code
 - Acceptable Use of Technology Policy
 - Data Classification and Cybersecurity Policies

Use of AI does **not** change the classification or protection requirements of PAU data. All users remain responsible for safeguarding institutional, confidential, and regulated information.

Ownership of AI-Generated Content

Content created using AI in the course of employment, academic work, or sponsored research at PAU is considered institutional work product, unless otherwise defined by contract, grant, or academic policy. Ownership and usage rights are governed by PAU's Intellectual Property Policy.



Out of Scope

This policy does not apply to personal use of AI tools that:

- Does not involve PAU data, systems, or resources
- Is performed entirely outside of university-related activities

However, individuals are expected to uphold PAU's ethical standards when representing the university or referencing their affiliation.

Roles and Responsibilities

The effective implementation and governance of this policy require collaboration across academic, administrative, and technical functions. Key stakeholders include:

- **President's Cabinet:** Provides executive oversight, approves policy updates, and ensures alignment with institutional strategy and ethical standards.
- **Provost's Office:** Oversees academic application of AI, ensuring responsible use in teaching, learning, and research.
- **Chief Information Officer (CIO) / vCISO:** Leads AI governance, data protection, and cybersecurity compliance; maintains the approved AI tool inventory and risk register.
- **Information Security & IT Operations:** Evaluate AI tools for data security, privacy, and system compatibility.
- **Legal Counsel and Compliance:** Review of vendor agreements, privacy practices, and regulatory compliance (FERPA, HIPAA, GLBA).
- **Human Resources:** Oversees employee training, ethics, and equitable AI use in staff development and hiring.
- **Institutional Review Board (IRB):** Ensures ethical standards for AI use in research, particularly involving human subjects.

Responsible Use of AI in Research & Administration

The use of Artificial Intelligence (AI) in research and administrative activities at Palo Alto University (PAU) must be conducted in a manner that ensures data protection, regulatory compliance, and ethical integrity.

All users are responsible for understanding the type of data being used and ensuring that AI tools are applied appropriately based on data sensitivity and risk.

Data Protection Requirements

The use of AI does **not** change the classification or protection requirements of PAU data.

- Restricted Data (FERPA, HIPAA, GLBA, PII, confidential research data) must not be entered into Public AI Tools under any circumstances
- Institutional data may only be used with Approved AI Tools, and only where appropriate safeguards are in place
- Users must apply data minimization principles (only provide the minimum necessary information to achieve the intended outcome)
- All users remain accountable for protecting data, regardless of the AI tool used

AI Usage by Data Classification

This reference shows which AI tools may be used with each of PAU's three data classifications — **Public**, **Private**, and **Restricted** — as defined in the Data Classification Policy. "Public AI Tools" refers to public, self-paid platforms with no PAU vendor agreement (e.g., ChatGPT, standalone Claude, Perplexity). "Approved AI Tools" refers to institutionally provided platforms operating under a PAU vendor agreement and access controls (e.g., Microsoft Copilot, Google Gemini, Claude Team).

Data Type	Public AI Tools	Approved AI Tools
Public / non-sensitive data Examples: • Press releases • Directory information (If not blocked by FERPA) • Course catalogs • Application and request forms • Campus maps • Event announcements • Course details.	 Allowed	 Allowed
Private (non-sensitive operational data) Examples: • University budget details • Procurement information • Research proposals • Internal audit reports	 Prohibited	 Allowed

Internal policies, procedures and guidelines		
Restricted Data (FERPA, HIPAA, GLBA, PII) Examples: - Protected Health Information (PHI) subject to the Health Insurance Portability and Accountability Act (HIPAA). - Student & education records (FERPA) like grades, transcripts, disciplinary records, etc. - Student personally identifiable financial information (GLBA) like income and tax return information, bank account numbers, student loan account numbers, balances, etc. - Financial account numbers covered by the Payment Card Industry Data Security Standard (PCI-DSS), which controls how credit card information is accepted, used, and stored.	✗ Prohibited	⚠ Conditional (requires approval and safeguards)

Risk-Based Use of AI

AI use at PAU should be evaluated based on risk and impact:

- Low Risk**
 - Drafting communications
 - Summarizing public or non-sensitive content
 - Brainstorming ideas
- Moderate Risk**
 - Processing internal documents
 - Supporting administrative workflows
 - Assisting with research preparation
- High Risk**
 - Handling sensitive or regulated data
 - Supporting decisions that impact students, employees, or research outcomes (including clinical data such as session transcripts, and decisions affecting clients served through PAU-operated clinics)
 - Automating processes with compliance, legal, or ethical implications

High-risk use cases require review and approval by IT/Security and/or appropriate governance bodies (e.g., IRB, Legal, Provost's Office).

Commented [DS4]: add clinical data including session transcripts; Supporting decisions that impact clients within clinics@PAU

Research Use Requirements

- All AI-assisted research must comply with PAU's Data Classification and Cybersecurity Policies
- Research involving human subjects must be reviewed and approved by the Institutional Review Board (IRB)
- AI may not be used to fabricate, falsify, or misrepresent research data or findings
- Researchers must ensure transparency and reproducibility when AI tools are used in research processes

Administrative and Operational Use

Departments seeking to use AI in administrative or operational workflows must ensure:

- **Risk Assessment** – Evaluate impact on data security, compliance, and accuracy
- **Data Protection** – Confirm appropriate handling of institutional and sensitive data
- **Human Oversight** – Ensure AI outputs are reviewed before use in decision-making
- **Transparency** – Disclose AI use where it materially impacts outcomes or communications

Clinical and Supervisory Use of AI

AI use within PAU's clinical training programs and affiliated training clinics raises considerations beyond those addressed elsewhere in this policy, given the sensitivity of client health information and the supervisory relationships involved in clinical training. The following requirements apply in addition to, not in place of, the rest of this policy:

- Clinical and health-related data — including therapy session transcripts, progress notes, chats with conversational AI agents used in supervision, and telehealth consent records — are Sensitive Data and Institutional Data and must be protected accordingly (see Appendix A).
- Written, informed client consent is required before AI tools are used in connection with client care, telehealth sessions, or clinical supervision. This is a higher bar than the general transparency requirement described in Ethical Principles for AI Use.
- Algorithmic bias in AI tools used for clinical purposes must be evaluated before the tool is adopted and monitored on an ongoing basis throughout its use.
- AI-generated clinical documentation (e.g., draft progress notes) must be reviewed by a licensed supervisor or other qualified staff member before it is relied upon or entered into the clinical record. AI outputs do not replace clinical judgment, including in identifying risk factors such as suicidality.
- Student trainees who use university-approved clinical AI tools under formal supervision remain responsible for reviewing and validating AI-assisted outputs. Ultimate clinical

Palo Alto University

responsibility rests with the supervising licensed clinician, consistent with PAU's clinical supervision policies.

- Questions about the use of AI in clinical settings should be directed to clinic leadership and the Provost's Office, in consultation with IT/Security and Legal Counsel.

Prohibited Activities

The following activities are strictly prohibited:

- Entering sensitive or restricted data into Public AI Tools
- Using AI to generate false, misleading, or fraudulent information
- Relying solely on AI outputs for decisions without human review
- Using AI in a manner that violates laws, regulations, contracts, or PAU policies
- Diagnosing or determining treatment
- Determining suicide, violence, or abuse risk
- Making hospitalization or mandated-reporting decisions

Reporting & Enforcement

Palo Alto University (PAU) is committed to ensuring the responsible and secure use of Artificial Intelligence (AI). All members of the PAU community share responsibility for complying with this policy and for reporting suspected violations.

Reporting Suspected Violations

Any individual who becomes aware of a potential violation of this policy is expected to report it promptly through the appropriate channel:

- **Students:** Report to the instructor or through the Academic Honor Code process
- **Faculty:** Report to the Provost's Office
- **Staff:** Report to their supervisor, Human Resources, or Information Technology
- **Security or Data Concerns:** Report directly to Information Technology / Information Security

Reports may include, but are not limited to:

- Unauthorized use of AI tools with sensitive or restricted data
- Misuse of AI resulting in academic dishonesty or research misconduct
- Use of unapproved AI tools in institutional processes involving PAU data
- Any activity that violates applicable laws, regulations, or PAU policies

PAU may provide additional mechanisms for confidential or anonymous reporting, where appropriate.

Review and Investigation

All reported concerns will be reviewed and, where appropriate, investigated in coordination with relevant university offices, which may include:

- Information Technology / Information Security
- Provost's Office
- People Operations
- Outside Legal Counsel, where applicable
- Institutional Review Board (IRB), where applicable

Investigations will be conducted in a manner consistent with existing PAU policies and procedures.

Severity and Response

Violations of this policy will be assessed based on severity, impact, and intent. Factors considered may include:

- The type and sensitivity of data involved
- Whether the violation was intentional or unintentional
- The potential or actual impact to individuals, the institution, or compliance obligations

Examples of severity levels may include:

- **Minor Violations**
 - Unintentional misuse of AI without exposure of sensitive data
 - Failure to properly cite AI-generated content
- **Moderate Violations**
 - Use of unapproved AI tools for institutional work involving internal data
 - Repeated non-compliance with policy requirements
- **Severe Violations**
 - Exposure of sensitive or regulated data (FERPA, HIPAA, GLBA)
 - Academic misconduct, research falsification, or fraudulent use of AI
 - Intentional violation of security, legal, or compliance requirements

Enforcement Actions

Enforcement actions will be handled in accordance with existing PAU policies and procedures:

- **Students:** Subject to the Academic Honor Code and disciplinary processes
- **Faculty:** Subject to review and action by the Provost's Office
- **Staff:** Subject to Human Resources policies and disciplinary procedures

Actions may include, but are not limited to:

- Required training or remediation
- Revocation of system or tool access

Palo Alto University

- Formal disciplinary action, up to and including termination or dismissal

Accountability

All members of the PAU community are accountable for:

- Understanding and complying with this policy
- Exercising appropriate judgment when using AI tools
- Always protecting institutional and sensitive data

Failure to comply with this policy may result in disciplinary action and, where applicable, legal or regulatory consequences.

Ethical Principles for AI Use

Artificial Intelligence (AI) at Palo Alto University (PAU) must be used in a manner that reflects the institution's commitment to integrity, accountability, and responsible innovation. The following principles guide all academic, research, and administrative use of AI:

- **Transparency**
Users must disclose the use of AI where it materially influences academic work, research, or institutional decision-making.
- **Accountability**
Human judgment remains responsible for all outcomes. AI-generated outputs must be reviewed, validated, and approved by the user prior to use.
- **Fairness and Equity**
AI must be used in a manner that avoids bias, discrimination, or inequitable outcomes, particularly in decisions affecting students, employees, or research subjects. For AI tools used in clinical contexts, algorithmic bias must be evaluated before adoption and monitored throughout ongoing use.
- **Integrity**
AI must not be used to misrepresent work, fabricate information, or undermine academic, research, or professional standards.

Commented [DSS]: Maybe add more for clinical decision making: Evaluate algorithmic bias before clinical tools are used/adopted and throughout use

AI Subscription and Licensing

Palo Alto University (PAU) manages the use of Artificial Intelligence (AI) tools through a controlled and risk-based approach to ensure data protection, regulatory compliance, and responsible use.

Institutional Approval of AI Tools

AI tools used for university-related activities must be reviewed and approved prior to use with PAU data or in institutional workflows.

- Only **Approved AI Tools** may be used to process, store, or generate content involving PAU institutional or internal data

Printed copies of this document are uncontrolled. In the case of a conflict between printed and electronic versions of this document, the controlled version published online prevails.

Palo Alto University

- Approval is granted by Information Technology and Security based on data protection, privacy, security, and compliance requirements
- All approved tools will be maintained in an **institutional AI tool inventory or register**

Use of AI tools that have not been approved is considered **unauthorized** when used with PAU data or for institutional purposes.

Use of Public and Individually Licensed AI Tools

Members of the PAU community may choose to use publicly available or individually licensed AI tools (e.g., ChatGPT, Claude, Gemini) at their own expense, subject to the following conditions:

- These tools must be treated as **Public AI Tools**
- **Sensitive or restricted data must not be entered into these tools under any circumstances**
- Use must comply with all PAU policies, including Acceptable Use, Data Management, and Information Security
- Users assume responsibility for any risks associated with the use of non-approved tools

Public or individually licensed AI tools may be used for **low-risk activities only**, such as:

- Drafting general communications
- Brainstorming or ideation
- Summarizing publicly available information

Prohibited Use of Non-Approved Tools

Non-approved AI tools may not be used for:

- Processing institutional, confidential, or regulated data
- Supporting administrative or operational workflows
- Making or informing decisions that impact students, employees, or university operations

Future Institutional AI Capabilities

PAU may evaluate, procure, and implement enterprise AI solutions that meet institutional requirements for security, compliance, accessibility, and integration.

As institutional AI tools become available:

- Users will be expected to utilize approved platforms for university-related work
- Additional guidance and standards will be provided to support adoption

Compliance and Accountability

All users are responsible for:

Printed copies of this document are uncontrolled. In the case of a conflict between printed and electronic versions of this document, the controlled version published online prevails.

Palo Alto University

- Using only approved AI tools when working with PAU data
- Understanding the distinction between Public and Approved AI tools
- Complying with this policy and all related data protection requirements

Failure to comply with this section may result in restricted access to systems, disciplinary action, or other enforcement measures in accordance with PAU policies.

Policy Review

This policy will be reviewed annually by the Provost's Office, CIO, Faculty representatives, and other stakeholders to ensure alignment with evolving AI technologies, ethical standards, and compliance requirements.

Commented [DS6]: Maybe there should be a Clinical AI section or addendum? Right now, clinical AI would be decided by IT and legal. Clinical representation would also be important.

I would also suggest clarifying student trainee liability when using university approved clinical AI tools under formal supervision

Appendix A: Definitions & Key Terms

For the purposes of this policy, the following definitions apply:

- **Artificial Intelligence (AI):** Systems or technologies that perform tasks typically requiring human intelligence, including reasoning, learning, problem-solving, language processing, and decision-making.
- **Generative AI:** AI technologies that create or generate new content—such as text, images, audio, video, or code—based on user input (e.g., ChatGPT, Microsoft Copilot, Gemini, Claude, DALL·E).
- **AI Tool:** Any software, platform, or embedded system that uses artificial intelligence to process data, generate content, or support decision-making. This includes both standalone tools and AI capabilities integrated into enterprise systems.
- **Approved AI Tool:** An AI tool that has been reviewed and authorized for institutional use by PAU Information Technology and Security based on data protection, privacy, security, and compliance requirements.
- **Public AI Tool:** An AI tool that is publicly available and not specifically approved or managed by PAU (e.g., free or individually licensed versions of generative AI platforms). These tools typically process data outside of PAU’s controlled environment.
- **Prompt:** Any input, instruction, or query provided to an AI system to generate a response or output.
- **AI-Generated Content:** Any content that is created, summarized, translated, or significantly influenced by an AI tool, including text, images, code, or data outputs.
- **Sensitive Data (Restricted or Protected Data)** - Information that is subject to legal, regulatory, or institutional protection requirements, including but not limited to:
 - Student education records (FERPA)
 - Health or clinical information (HIPAA)
 - Financial or financial aid data (GLBA)
 - Personally identifiable information (PII)
 - Confidential institutional or research data
- **Institutional Data:** Any data created, collected, stored, or managed by PAU during its academic, research, clinical, or administrative operations, regardless of format or location. This includes HIPAA-protected clinical data, such as therapy session transcripts, chats with conversational AI agents used in supervision, and signed telehealth or AI-supported supervision consents generated through PAU-operated training clinics.
- **AI-Assisted Decision-Making:** The use of AI outputs to inform, support, or automate decisions in academic, research, or administrative contexts. All such decisions require appropriate human review and accountability.

Commented [DS7]: An area to expand on is clinical operations. Is HIPAA protected data such as session transcripts, chats with a conversational AI agent in the supervisor platform, or signed consents for telehealth and AI supported supervision part of Institutional Data?

Appendix B: Resources & References

The following resources are provided to support responsible and effective use of Artificial Intelligence (AI) at Palo Alto University (PAU). All faculty, staff, and students are encouraged to consult these references for additional guidance.

PAU Internal Policies & Guidelines

- Academic Honor Code –
- Faculty Handbook – Available through the Provost’s Office Portal
- Acceptable Use of Technology Policy –
- Data Classification & Cybersecurity Policy –
- Intellectual Property Policy –

Federal Compliance & Regulatory Resources

- FERPA Guidance: <https://studentprivacy.ed.gov/ferpa>
- HIPAA Guidance: <https://www.hhs.gov/hipaa/>
- GLBA Safeguards Rule: <https://www.ftc.gov/business-guidance/privacy-security/gramm-leach-bliley-act>
- IRB Guidelines: <https://www.hhs.gov/ohrp/regulations-and-policy/regulations/common-rule/index.html>

Technology & Tool Guidance

- List of PAU Approved AI Tools – Maintained by PAU IT (IT Support Portal) [Insert direct link to IT Support Portal AI tool inventory]
- Request Form for New AI Tools –
- AI Detection Tools – Faculty advised to consult the Provost’s Office before adoption. Faculty are advised that AI detection tools can produce false positives and must not be used as the sole basis for initiating an Academic Honor Code violation.

Reporting & Support Contacts

- Provost’s Office – Academic policy support
- IT Support: support@paloaltou.edu
- Institutional Review Board (IRB): irb@paloaltou.edu
- Human Resources: hr@paloaltou.edu

Commented [GS8]: Can we provide a link to this or tell us exactly where to find them?

Commented [GS9]: What do you think about adding this? Faculty are advised that AI detectors exhibit false positives and must not be used as the sole basis for initiating an Academic Honor Code violation.

FAQ:

1. **What is considered Artificial Intelligence (AI) under this policy?**
AI includes systems or tools that perform tasks requiring human-like intelligence, such as text generation, image creation, data analysis, or automated decision-making. Examples include ChatGPT, Microsoft Copilot, Gemini, and similar technologies.
2. **Does this policy apply to all AI tools, even free ones?**
Yes. The policy applies to all AI tools—whether free, paid, or provided through institutional licenses—when used for academic, research, or administrative purposes.
3. **Can students use AI for coursework?**
Students may use AI for ideation, translation, or research assistance only if explicitly permitted by the instructor. Unauthorized or uncredited AI use is treated as academic dishonesty under PAU's Academic Honor Code.
4. **How should AI-generated content be cited?**
AI outputs must be acknowledged in the same way as any external source. Students should include the tool name, version (if available), and prompt or description of use (e.g., "Generated using ChatGPT, prompt: summarize Chapter 3 of...").
5. **Are faculty required to allow AI use in their courses?**
No. Faculty have full discretion to determine if and how AI tools may be used. These expectations should be clearly stated in the syllabus (e.g., permitted, guided, or prohibited use).
6. **What precautions must be taken when using AI for research?**
Researchers must not input any restricted or protected data (e.g., FERPA, HIPAA, GLBA) into public AI tools. All AI-assisted research must comply with PAU's Data Management and Cybersecurity Policies, and may require Institutional Review Board (IRB) approval.
7. **Can staff use AI for administrative tasks?**
Yes, when approved and aligned with PAU data protection standards. Examples include drafting communications or improving workflows. Departments should conduct Risk and Equity Assessments before operationalizing AI use.
8. **Can AI replace staff or faculty positions?**
No. The replacement of human roles with AI will not occur without ethical review, stakeholder consultation, and approval from PAU leadership.
9. **What are some examples of prohibited AI use?**
 - Entering confidential or student data into external AI systems
 - Generating discriminatory or harassing content
 - Using AI to falsify information or credentials

Printed copies of this document are uncontrolled. In the case of a conflict between printed and electronic versions of this document, the controlled version published online prevails.

- Violating laws, contracts, or institutional policies through AI-generated actions

10. Who do I contact with questions about using AI tools?

- **Provost's Office:** Academic and faculty-related AI guidance
- **Information Technology (IT):** Tool approval, risk assessments, technical support
- **Human Resources:** Staff-related use and compliance inquiries
- **Institutional Review Board (IRB):** Research ethics and approvals

11. What happens if someone violates this policy?

Violations are handled according to existing PAU procedures:

- **Students:** Academic Honor Code process
- **Faculty:** Reported to the Provost's Office
- **Staff:** Reported to the supervisor or Human Resources

12. How often will this policy be reviewed?

The AI Policy will be reviewed annually, consistent with the Policy Review section of this document by the Provost's Office, CIO/vCISO, and faculty representatives to ensure it remains aligned with evolving technologies and regulations.

Commented [DS10]: says one year on page 11