

California State University, Northridge (CSUN)

AI Principles and Guidelines

Working Draft — 2026-04-17

Introduction

Artificial intelligence tools are increasingly present in higher education, affecting how universities teach, learn, conduct research, and operate. It will also be an inevitable part of our students' futures. The purpose of these guidelines is to support responsible AI use by those who choose to use it, provide guardrails that protect the university community, and respect the professional judgment of faculty and staff who determine that AI is not appropriate for their work. Uses will vary by discipline and operational need, and these guidelines are designed to accommodate that variation.

This document was developed under the CSUN AI Task Force's AY25/26 charge to help the campus adopt AI in ways that are ethical, effective, and aligned with CSUN's academic mission. It is **guidance, not binding policy, and will be revisited as AI and campus needs evolve**. It is being developed in conjunction and collaboration with other committees and units at CSUN. This document will link to documents produced by those other units when ready.

Universal Principles

The following principles apply to all members of the CSUN community (faculty, students, staff, and administrators) across all domains. They are anchored in the six CSU systemwide AI Guiding Principles (csun.edu/ai/csu-ai-guiding-principles).

Integrity and Human Accountability. AI is a tool that supports human judgment, not a replacement for it. Anyone using AI tools remains responsible for the accuracy, completeness, and appropriateness of the decisions and outputs that result. For consequential decisions affecting students, employees, or research subjects, a qualified person should meaningfully review the AI's output before it is acted upon. That person should have responsibility for the function, appropriate subject knowledge, and authority to approve the result.

Privacy and Data Protection. AI use at CSUN follows CSU data governance policy, applicable law, and CSUN security standards. Under CSU data classification, Level 1 is the most sensitive category (including Social Security numbers and other protected personal data) and Level 2 covers internal data restricted from public disclosure. Because of legal and contractual obligations, Level 1 data must not be entered into any non-approved AI system, and Level 1 and Level 2 data may only be processed in campus-approved tools. Information Technology maintains the official approved tools list and can advise on enterprise products not yet on it, some of which can be configured to meet CSUN requirements.

Transparency and Disclosure. Authors should first check the AI-disclosure policies of the intended recipients of their work (journals, funding agencies, course instructors, or

administrative offices) and follow them. Within CSUN, community members should be open about when and how AI was used, and should document it when AI played a meaningful role in the work.

Equity, Access, Inclusion, and Justice. AI systems can produce errors, reflect the biases of their training data, and amplify systematic discrimination. Users should therefore critically evaluate AI outputs for bias and consider unintended impacts on individuals and marginalized groups; tools that fail this test are not appropriate as educational resources and not reliable as research methods. When AI tools are required in a course or service, the university should ensure equitable access, and AI-supported services should meet campus accessibility expectations rather than creating barriers for individuals with disabilities. For accessibility assistance, contact udc@csun.edu.

Learning, Development, and Risk Management. AI should enhance understanding, professional readiness, and operational efficiency, not bypass them; critical thinking and independent judgment remain core competencies regardless of how capable AI tools become. Departments and offices adopting AI should also think through real operational considerations in advance, including that AI can produce incorrect outputs, introduce automation errors, and create service disruptions.

Domain 1: Teaching and Learning

Course-Level AI Policies

Faculty decide how and whether AI supports their learning objectives. CSUN offers three suggested policy templates that faculty can communicate in their syllabus and assignment instructions. Where AI use is permitted under any of these templates, students should disclose and cite it.

- **AI-Restricted.** AI tools are not used for course components unless the instructor specifies otherwise. Potential use: courses where unassisted demonstration of competency is central (e.g., introductory writing, formal proofs, clinical documentation, world language acquisition).
- **AI-Assisted.** AI tools are permitted for specific tasks as described in the instructions. Potential use: courses integrating AI as a productivity or research tool alongside disciplinary learning objectives.
- **AI-Integrated.** AI tools are expected as part of coursework so students develop AI literacy alongside domain skills. Potential use: courses where AI competency is a learning outcome (e.g., data analysis, design, strategic planning).

When a course does not state an explicit AI policy, students should ask their instructor before using AI on graded work. Students who wish to opt out of AI use for ethical or religious reasons should also discuss with their instructor whether optional assignments are available.

Assessment and Academic Integrity

When AI is used in coursework, faculty should assess students on their learning process, not just their outputs, to help build AI literacy. This can include drafts and process artifacts, reflection on reasoning and choices, in-class work, scaffolded checkpoints, or oral explanation of the work.

Students remain accountable for the accuracy and integrity of all AI-assisted work. Concerns of misuse of AI should be referred to the Office of the Vice President of Student Affairs, as reflected in the CSUN Academic Integrity Policy. Some examples of misuse of AI include: submitting AI-generated work as one's own unassisted work in a course where AI use has not been permitted, and submitting AI-generated citations that do not correspond to real, verifiable sources (whether or not the student was aware they were fabricated).

Faculty who use AI in grading or feedback, or AI detectors, should disclose this to students and define boundaries in their course policy. AI can support grading but should not replace the instructor's professional judgment, and faculty should review and edit AI-assisted evaluations before releasing them to students. AI detection tools may be helpful as one input, but should not be the sole basis for an academic integrity outcome; documented false-positive rates disproportionately affect non-native English writers and students with certain learning disabilities.

AI Literacy and Course Materials

Programs should integrate AI literacy into their curricula, with an emphasis on critical evaluation of outputs, understanding of bias, and ethical judgment, to prepare students for disciplines and workplaces where AI is increasingly present. Faculty should also help students think about what materials are appropriate to share with AI tools in the first place. Students should not load instructor-authored materials or course-licensed primary sources into non-CSUN-approved AI tools without the instructor's permission.

Domain 2: Research

This domain applies to all CSUN research activity: faculty research, student research (graduate and undergraduate), and staff research in support roles.

Disclosure and Attribution

AI cannot take accountability, respond to peer review, or attest to findings. For these reasons, AI should not be listed as an author. When submitting manuscripts, proposals, or theses, researchers should check the AI-disclosure policies of the target journal, funding agency, or other outlet and follow them; because these policies vary and continue to evolve, researchers should check regularly. Where an outlet lacks a stated policy, researchers should disclose substantive AI involvement in the methods or acknowledgments section as a default.

Data Security and Technical Risks

Sensitive, regulated, or human-subject data should not be entered into non-CSUN-approved AI systems. Researchers working with human subjects should consult the CSUN Institutional Review Board about whether planned AI use is covered by existing protocols or requires amendment. As a best practice, researchers should anonymize data by stripping identifiers before using any AI tool. Using AI to fabricate, falsify, or manipulate data or results is research misconduct, the same as it would be with any other tool.

Because AI systems can generate hallucinated content (fabricated facts, nonexistent citations), researchers should independently verify AI-generated material before relying on it.

AI tools can also be tricked by misleading information in documents or websites into ignoring rules, producing biased results, or taking unauthorized actions. Researchers should build safeguards into automated workflows and should evaluate whether the AI models they use may carry training-data biases into their findings.

Infrastructure

IT provides colocation support for grant-funded AI research infrastructure and partners with faculty to enable access to external AI research computing resources available through federal programs and other providers. IT can also advise on appropriate security, compliance, and architectural considerations, and assist with access pathways and integration support where applicable. This support model will be revisited periodically to align with evolving researcher needs and available funding mechanisms.

Domain 3: Administration and Operations

Use Cases

In administrative and operational contexts, staff should use AI primarily as a support for routine, drafting, and early-stage work that will be reviewed before it goes out. Common examples include drafting routine communications, analyzing datasets for reporting, handling first-level customer service with clear escalation paths, automating repetitive tasks through approved tools, and brainstorming.

Limits and Escalation

Several categories of work sit outside what AI should be used for in an administrative context. Level 1 and Level 2 data must not be entered into unapproved tools. Final decisions in hiring, performance evaluation, disciplinary processes, or other actions with legal or contractual consequences are not appropriate uses of AI; neither are automated systems that determine access to accommodations, employment, financial aid, or benefits without human review. New AI applications that will interface with campus data, enterprise systems, or processes should be brought to Information Technology for security, procurement, and governance consultation before deployment.

Oversight and Disclosure

A qualified employee should review AI-assisted output before release. Departments and offices should also acknowledge AI assistance when it contributes significantly to a document or decision, for example, when AI generates large portions of content, summarizes material for approvals, or provides recommendations. Minor uses such as grammar or spelling assistance do not require disclosure.

Procurement, Training, and Support

New AI tools go through campus procurement review, which covers vendor security, data protection, and accessibility.

Staff AI training is available at www.csun.edu/ai/training.

For compliance and security questions, contact Information Security (iso@csun.edu). For general AI support on authorized campus AI tools and resources, contact helpcenter@csun.edu.

AI disclosure. *AI was used for initial drafting and later restructuring, editing, and tightening of this document. All output was reviewed (and continues to be reviewed) by members of the CSUN AI Task Force, who made all substantive decisions on content and policy positions.*