Principles and Guidance for the Responsible Use of Artificial Intelligence

The Office of the Provost and the AI Steering Committee | Effective: [Date] | Next Review: [Date]

AISteering@DU.EDU

As Artificial Intelligence (AI) reshapes how we teach, learn, and work, DU is committed to responsible use of AI on our campus that reflects our deepest institutional values. These principles and guidelines define our approach and hold us accountable to our community.

The University of Denver Artificial Intelligence Steering Committee (DU AISC) is established as a working group dedicated to the governance, evaluation, approval and support of artificial intelligence (AI) technologies for operational use across University departments.

Table of Contents

Part I: Core Values
Part II: Ethics Statements
Part III: General Guidelines
Part IV: Guidelines by Domain
- A. Teaching & Learning
- B. Research
- C. Operations
- D. Promotion & Tenure

Part V: Use Cases
Part VI: Frequently Asked Questions

DRAFT

Part I: Core Values

Human Judgment Remains Central

Artificial Intelligence is a powerful tool, but it does not replace the expertise, care, and human dignity that define our scholarly community. At DU, we believe every AI-assisted decision (e.g., research, instruction, service, advising, operations, administration) must remain subject to informed human judgment. AI supports the essential work of our faculty, staff, and students; it does not direct or replace it.

Oversight as Scholarly Responsibility

Just as we hold ourselves to standards of academic rigor in research and peer review, we hold ourselves to the same standards in AI use. Deploying an AI system without meaningful human review is inconsistent with our commitment to accuracy, fairness, and intellectual honesty. We do not accept outputs simply because a system produced them.

Transparency, Privacy, and Accountability

Any use of AI in our institutional work must be visible and explainable to students, faculty, staff, and the broader public we serve. We should disclose when AI has shaped a consequential decision or recommendation, and a human being must remain accountable for the outcome. We are equally committed to protecting the privacy of personal and sensitive information in every context, including those in which AI has been used.

Equity, Accessibility, and Active Vigilance

AI systems reflect the data and design choices that formed them, and those decisions can embed bias. We are committed to providing access so all members of our campus community, students, faculty, and staff, can benefit from AI regardless of background, identity, disability, financial circumstances, field of study, or prior technical experience. We assess AI tools for disparate impacts on marginalized communities and act when we find inequities, triggering a formal review, honest reporting, and a defined remediation plan.

Part II: Ethics Statements

On the Integrity of Learning

DU's educational mission depends on authentic intellectual development. AI use in student work must be governed by guidelines that preserve the learning process itself, not just the appearance of a finished product. As a subject of study, we support the purposeful integration of AI. As a learning aid, we are committed to building AI literacy across disciplines so that every student can critically evaluate AI-produced content, recognize bias and misinformation, and use AI tools in ways consistent with course expectations and University values.

On the Environmental and Societal Costs of AI

The use of AI carries environmental consequences beyond our campus, including the energy and water demands of AI infrastructure. In keeping with our commitment to sustainability, DU will consider environmental impact alongside other ethical criteria when evaluating AI enterprise tools available to the DU community and will seek to avoid unnecessary or wasteful use of AI systems where simpler approaches suffice.

On the Limits of Automation

DRAFT

DU will not use AI to make final determinations in consequential matters, including admissions, grading, discipline, hiring, or resource allocation, without meaningful human deliberation. Automated efficiency ought never to come at the cost of individual dignity or due process.

On Verification and Review

DU faculty, staff, administrators, and students who use AI tools bear responsibility for verifying the outputs those tools produce. Accepting AI-generated content, analysis, or recommendations without critical review is a failure of professional duty.

On Avoiding Automation Bias

At DU, we commit to training members of our community to recognize automation bias, the tendency to over-rely on AI outputs because they appear authoritative or objective. Critical engagement with AI is not optional. We will provide the education, professional development, and support needed to sustain genuinely independent human review.

On Continuous Institutional Review

DU reviews and scrutinizes all AI systems continuously. Continuous oversight is essential to responsible AI governance. Before and during their use, DU will assess AI tools for accessibility, privacy, security, reliability, and risk of disparate impacts. We commit to regular, documented review; to engaging students and campus partners in that process; and to retiring or replacing tools that fall short of our institutional values.

Part III: General Guidelines

These guidelines apply to all faculty, staff, and students using AI in any DU activity.

1. Work Within University Policies

Members of the University community are responsible for using AI tools in a manner consistent with existing University policies, including the Technology Acquisition Policy (IT 13.10.040), Data Classification Policy (IT 13.10.051), Institutional Data Management Policy (IT 13.10.050), and Acceptable Use Policy (IT 13.10.010).

2. Protect Confidential, Personal, and Restricted Information

Entering information into an AI tool is similar to posting it on a public website, noting that it may become part of training data and be reproduced for others. Data classifications are specified in IT 13.10.051 Data Classification. Internal, Confidential, and Restricted Data (i.e., financial information, personally identifiable information, health information, university-specific information) may only be used with AI tools when: (a) the data is classified as public, or (b) the tool has been reviewed and approved by the appropriate Data Governance Committee as required by the Technology Acquisition Policy with appropriate data protection terms in place.

3. You Are Responsible for What You Use

AI-created content may be biased, inaccurate, fabricated, or contain copyright-protected material. Before using any AI output — in a document, decision, communication, or submission — you must review it for accuracy, appropriateness, and compliance. The person who uses AI remains responsible for the final work product.

4. Disclose and Attribute AI Use

DRAFT

Disclose any AI use that materially adds to a work product, submission, publication, or decision. This includes use in drafting, analysis, translation, summarization, image generation, or coding. Follow the attribution standards of your discipline, the relevant journal or publisher, and any applicable University policy. For example, if a staff member uses AI to analyze survey responses that inform a recommendation to leadership, the recommendation memo states that AI was used to categorize and summarize the open-ended responses.

5. Follow Approval Requirements

Any University employee who desires to utilize an AI tool that process Internal, Confidential, or Restricted Research Data, or that may affect individuals' legal rights or physical safety, must submit such AI Tool for review to the Technology & Data Services Review Committee (which may elevate the review to the AI Steering Committee) and obtain Committee approval before deployment.

6. When in Doubt, Ask

AI technology evolves rapidly, and no set of guidelines can anticipate every situation. Contact the IT Help Center (du.edu/it) to connect with the right people. The DU AISC welcomes questions and use-case consultations.

Part IV: Guidelines by Domain

A. Teaching and Learning

A.1. For Faculty

A.1.a. Setting Expectations

Faculty bear primary responsibility for establishing clear AI use expectations in every course. Faculty may set different AI rules by course or assignment to fit their discipline or learning objective. Communicate expectations in the syllabus at the start of each term and specify which AI tools (if any) are permitted, for which tasks, and what disclosure students are expected to provide. Faculty are encouraged to model responsible AI use by disclosing their own use of AI in course design and/or management.

Faculty are encouraged to review the Office of Teaching and Learning's resources on AI in the Classroom.

A.1.b. Permissible Instructional Uses

Faculty may use AI tools to support course design and material development; developing rubrics; remediating inaccessible course materials; recommendations for adapting course work for accessibility accommodations; designing presentation materials; and drafting administrative materials (emails). Faculty are encouraged to review all AI-generated content before using it and remain responsible for accuracy and quality.

A.1.c. Prohibited Uses

Faculty may not use AI tools to: assign final grades without human review; input identifiable student information into non-enterprise tools (e.g. tool that have not been formally adopted by DU and may not be FERPA compliant); generate student records, recommendations, or disciplinary findings without independent verification; or substitute AI evaluation for substantive engagement with student work.

A.1.d. Designing for AI Literacy

DRAFT

Faculty are encouraged to design assessments and learning activities that build students' capacity to evaluate, critique, and use AI responsibly. DU's Office of Teaching and Learning and Center for Ethical AI are resources for curriculum development support.

A.2. For Students

A.2.a. Governing Rule

Students are expected to abide by the Honor Code and by the AI-specific expectations established by each instructor. When an instructor has not addressed AI use, students should assume it is not permitted and ask for clarification before using it.

B. Research

DU researchers who receive external or federal funding, particularly from NSF and NIH, are subject to those funder or agency specific AI requirements. NSF and NIH serve as the University's primary guidance for federally-supported work. Because these policies continue to evolve, researchers should consult the current requirements directly:

- NSF AI Policy: nsf.gov/policies/ai
- NIH AI Guidance: osp.od.nih.gov/policies/artificial-intelligence/

Other general guidance for research includes the following considerations.

B.1. Data Protection and Confidentiality

Research data entered in an external (to DU) AI tool may be transmitted to and retained by outside providers, creating significant risk of unintentional disclosure. Researchers must not use external AI tools with any of the following without specific tool approval from the AISC:

- Data subject to IRB confidentiality protections or Certificates of Confidentiality
- Identifiable human subjects data, including HIPAA-protected health information
- Genomic data or data from communities with special sharing provisions (e.g., American Indian/Alaska Native participants)
- Restricted Research Data, export-controlled data, or data under non-disclosure agreements

Contact the Office of Research and Sponsored Programs before using AI tools with any research data whose classification is uncertain.

B.2. Peer Review

Peer review depends on confidentiality. Uploading unpublished manuscripts, grant applications, or proposals to an AI tool can expose that material to third parties and violate the trust of the authors. For this reason, DU aligns its guidance with the standards set by the National Institutes of Health (NIH) among other standards. For example, NIH prohibits peer reviewers from using generative AI tools, including large language models, to analyze or formulate peer review critiques. Using AI in peer review constitutes a breach of confidentiality under NIH policy. Researchers serving as reviewers for NIH or other funding sources, or federal agencies should assume this prohibition applies unless the agency explicitly states otherwise.

B.3. Attribution and Publication Integrity

Faculty are encouraged to follow AI disclosure practices in publications, grant applications, and conference submissions in accordance with the standards of the relevant discipline, journal, funding agency, and publisher.

DRAFT

Authorship cannot be attributed to an AI system. Researchers are responsible for the accuracy, validity, and integrity of any AI-assisted work product.

B.4. Bias, Validity, and Biosecurity

Researchers using AI for data analysis, literature review, or content generation should validate outputs for accuracy, bias, and methodological soundness before relying on them. AI tools used in human subjects research that may affect participants' rights or welfare must be submitted to IRB for review. Research involving biological agents informed by AI must comply with DU's biosafety and dual-use research oversight policies consistent with NIH guidance.

B.5. Research on AI

Research that studies, tests, or evaluates AI systems is subject to all standard research integrity requirements. Researchers are encouraged to engage with DU's Center for Ethical AI (CEAI) to connect with the broader campus research community.

C. Operations and Administration

C.1. Decision-Making

AI tools may support but not replace human reasoning in decisions affecting individuals' employment, student status, financial aid, housing, discipline, or access to University services. Any AI system used in these contexts must be reviewed and approved by the DU AISC and must be regularly tested for bias and accuracy.

C.2. Communications and Content

Staff may use approved AI tools to draft internal communications, summarize documents, generate reports, and support administrative workflows. Staff must review all AI-assisted outputs before distribution. University communications should not be represented as wholly human-authored when AI materially contributed to them.

C.3. Procurement and Vendor Tools

When evaluating or procuring software that incorporates AI functionality, staff must follow the Technology Acquisition Policy (IT 13.10.040) and consult with IT to determine whether AISC review is required. Vendor contracts must include appropriate data protection terms. Treat AI functionality embedded in existing platforms as a new AI tool for review purposes if it processes University data.

C.4. Records and Privacy

Do not use AI tools to process student records, personnel files, health information, donor records, or other protected data without appropriate approval. Staff are reminded that FERPA, HIPAA, and other applicable laws govern data use regardless of the tool involved.

C.5. Meeting Notes

Staff may use approved AI tools to help capture, summarize, or transcribe meeting discussions. When AI is used to record or transcribe any meeting, attendees must be informed before the recording begins. Meeting notes and transcripts may contain sensitive University information; use only approved tools, and the responsible staff member must review AI-generated summaries for accuracy and appropriateness before distribution. AI-generated meeting records are University records and must be retained in accordance with DU's records retention schedules. Meetings that involve personnel matters, student issues, legal concerns, or other confidential topics require heightened care, and distribution must be limited to appropriate parties.

DRAFT

C.6. Performance Reviews

Performance evaluations involve consequential determinations about individuals' employment and professional standing and must reflect independent human deliberation. Supervisors may use approved AI tools for administrative tasks related to the review process, such as formatting documents or organizing prior feedback, but substantive evaluative content, including assessments of performance, development needs, and employment recommendations, must be authored by the responsible supervisor or manager and independently reviewed. AI must not generate performance ratings, written evaluations, or disciplinary findings without full human authorship and review. If AI tools materially contributed to any documents used in a review, the affected employee must be informed. Final accountability for all review outcomes rests with the supervisor or administrator conducting the review.

D. Promotion and Tenure

D.1. For Candidates

Faculty candidates may use AI tools to assist with drafting, organizing, or editing P&T materials (personal statements, teaching philosophies, research narratives, CVs), provided such use is disclosed in the materials submitted. AI use in research outputs submitted as evidence of scholarship must conform to the attribution and disclosure standards in Section IV-B. Candidates should consult with their department chair or dean's office regarding unit-specific expectations before submitting materials.

D.2. For Reviewers and Committees

AI tools must not be used to draft substantive evaluative text in Promotion & Tenure reviews, specifically assessments of candidate quality, recommendation language, or summary judgments, without disclosure and human review. AI may assist with administrative tasks such as formatting, scheduling, or summarizing publicly available materials. Final evaluative judgments must reflect independent human deliberation.

D.3. Evaluating AI-Assisted Scholarship

Committees evaluating scholarly work that involved AI assistance should focus on the intellectual contribution of the candidate, meaning the conceptual framing, methodological completeness, interpretive judgment, and scholarly significance, rather than penalizing appropriate and disclosed AI use. AI-assisted work that meets disciplinary standards and is properly attributed should be evaluated on its scholarly merits.

D.4. Note on Evolving Norms

Units are encouraged to develop and articulate unit-level expectations for AI in promotion & tenure and to communicate those expectations to candidates and reviewers in writing. The Provost's Office, in consultation with Faculty Senate, will issue additional guidance as norms and practices evolve.

Part V: Use Cases

The following examples illustrate how these guidelines apply in practice. This list is not exhaustive. As more use cases develop, **this content will likely move to the AI website.**

Teaching and Learning

DRAFT

Use Case	Status	Details Matter
A faculty member uses Microsoft Copilot to draft an initial course rubric, reviews and revises it before sharing with students.	**Permitted**	AI-assisted drafting with human review using a Tier 1 tool.
A faculty member uses an approved AI tool to generate quiz questions, reviews them for accuracy and bias, then incorporates them into a Canvas assessment.	**Permitted**	Human review required; no student data entered.
An instructor uses an approved AI tool to provide real-time feedback on student writing. The tool meets data protection requirements and students are informed prior to use.	**Permitted with Conditions**	Tool must be AI Steering Committee approved; students must be notified before use.
An instructor enters identifiable student assignment text into an unapproved commercial AI tool to generate feedback.	**Not Permitted**	Student data may not be entered into unapproved tools.

Research

Use Case	Status	Details Matter
A researcher uses an AI tool to conduct a preliminary literature search and identify relevant papers, then independently verifies all sources before citing them.	**Permitted**	AI assists discovery; researcher independently verifies all sources.
A faculty member uses AI to assist with drafting a grant proposal narrative, discloses use per funding agency requirements, and reviews all content for accuracy.	**Permitted with Conditions**	Disclosure required per funder and publisher standards.
A researcher enters de-identified survey response data into an unapproved commercial AI tool for thematic analysis.	**Not Permitted**	Data classification and tool approval must be confirmed before use. Contact ORSP if uncertain.
A research team deploys an AI system to screen job applicants for a funded project without prior AISC approval.	**Not Permitted**	AI systems affecting individuals' rights must be AISC-approved and tested for bias before deployment.

Operations and Administration

Use Case	Status	Details Matter
An AI Steering Committee uses an approved tool to identify leading practices for AI policy and guidelines in university settings.	**Permitted**	Consequential policy and operational guidance require independent human deliberation.

An HR staff member uses Microsoft Copilot to draft a job posting, reviews it, and posts it.	**Permitted**	Tier 1 tool; no confidential data; human review completed.
A department coordinator uses an approved AI tool to summarize notes from a public committee meeting.	**Permitted**	Public information; Tier 1 tool; summary reviewed before distribution.
A staff member uses an approved AI tool to transcribe a departmental meeting. Attendees are informed at the start of the meeting.	**Permitted with Conditions**	Approved tool required; attendees must be informed; notes reviewed before distribution; retained per records schedule.
A financial analyst uses an unapproved AI tool to analyze confidential budget projections.	**Not Permitted**	Confidential data may only be processed with AISC-approved tools.
An admissions counselor uses an AI system to rank applicants and sends acceptance decisions based solely on those rankings.	**Not Permitted**	Consequential decisions affecting individuals require independent human deliberation; AI may not be the sole decision-maker.

Promotion and Tenure

Use Case	Status	Details Matter
A faculty candidate uses AI to improve the clarity and organization of their teaching philosophy statement and discloses this use in submitted materials.	**Permitted with Conditions**	Disclosure of AI assistance required in all submitted P&T materials.
A supervisor uses AI to help organize notes and prior feedback before authoring a performance review, then writes the evaluative content independently.	**Permitted with Conditions**	AI may assist with organization; evaluative judgments must be independently authored and reviewed by the responsible supervisor.
A P&T committee member uses AI to draft the committee's written evaluation of a candidate without disclosure.	**Not Permitted**	Evaluative judgments must reflect human deliberation; any AI assistance must be disclosed.

Part VI: Frequently Asked Questions

(As more questions surface, this content will likely move to the AI website.)

Can I use ChatGPT or other commercial AI tools for my work at DU?

It depends on what information you're using and whether the tool has been approved. For tasks including only publicly available or non-sensitive information, you may use many commercial tools with care. For anything involving student data, personnel information, confidential research data, or other non-public University information, you must use a tool that has been reviewed and approved. When in doubt, contact IT (du.edu/it) before proceeding.

Do I have to disclose every use of AI?
You are expected to disclose AI use that materially contributes to a work product — a document, analysis, decision, or submission. Incidental uses (e.g., using AI to check grammar in an internal email) generally do not require formal disclosure, but transparency is always encouraged. In academic and research contexts, err on the side of disclosure.

What happens if I'm not sure whether a tool is approved?
Contact the IT Help Center or consult the approved tools list. Do not use an unapproved tool with University data while waiting for a determination. If you need to use a tool urgently, contact IT to discuss interim options.

I want to use an AI tool for my research that isn't on the approved list. What do I do?
Submit a request to the AI Steering Committee (AISC) through [link/email]. Describe the tool, the data you plan to use with it, and the research purpose. Do not begin using the tool with University data until approval is granted.

What are the rules for using AI in promotion and tenure materials?
Candidates may use AI to assist with drafting and editing P&T materials with disclosure. Committees must make their evaluative judgments independently and disclose any AI assistance. Units should clarify their specific expectations in writing. See Part IV-D for full guidance.

Who oversees AI at DU, and how can I raise a concern?
The AI Steering Committee (AISC) is responsible for oversight of University AI use. To raise a concern, report a tool that may be causing harm, or request a review, contact [email/link].

How often will these guidelines be updated?
The AI Steering Committee will evaluate these guidelines at least annually and will issue updates when significant changes in technology, law, or University practice warrant them. The AI Steering Committee will communicate changes to the campus community through email announcements and University newsletters.

DRAFT

For questions about this guidance, contact the AI Steering Committee at [email]. For technology-specific questions, contact the IT Help Center at du.edu/it.

AI helped create this document, specifically to audit its completeness and review contradictions or inconsistencies. The entire subcommittee has reviewed, edited, and approved the current version.