

UTEP Guidelines for the Responsible Use of AI in Research

Table of Contents

1.	Executive Summary.....	2
2.	Purpose	2
3.	Scope	2
4.	Relationship to Existing Policies and Governing Authorities.....	3
5.	Relationship to Academic Freedom	4
6.	Definitions.....	4
7.	Core Principles	6
7.1.	Institutional Risk-Tolerance Statement.....	6
8.	Permissible and Discouraged Uses	7
9.	Data Governance & Classification	9
9.1.	Data Use Agreements (DUAs)	10
10.	Human Subjects, Re-identification Risk, and HRPP Considerations	11
11.	AI Tool Approval, Procurement & Vendor Risk	12
12.	Export Controls & Controlled Unclassified Information (CUI)	12
13.	Suggested Unified Disclosure Language	13
	Venue-Specific Disclosure Requirements.....	14
14.	Compliance and Sponsor-Facing Guidance.....	14
14.1.	AI Use in Idea and Hypothesis Development	14
14.2.	AI Use in Proposal Development - Agency-Specific Positions.....	15
14.3.	PI Certification Responsibility	15
14.4.	AI Use During Research Conduct	16
15.	Records & Retention.....	16
16.	Privacy & Protected Health Information (PHI)	17
17.	Thesis and Dissertation Research	17
18.	Foreign Influence, Conflicts of Interest, and Research Security	18
19.	Training & Competency	19
20.	Governance & Questions.....	20
20.1.	Compliance and Consequences	20
21.	Intellectual Property	21
22.	Review Cycle.....	22
23.	Appendix A - References (Selected).....	23

1. Executive Summary

Over the past decade, artificial intelligence (AI) has moved from the periphery of the research enterprise to its center, reshaping how data is collected, analyzed, and interpreted, and how results are communicated and visualized. These guidelines respond to that shift. They establish University-wide standards for the responsible, secure, and ethical use of Artificial Intelligence (AI) including Generative AI (GenAI) in research conducted at The University of Texas at El Paso (UTEP) and apply to all individuals engaged in UTEP research activities, including faculty, staff, students, postdoctoral researchers, and contractors regardless of funding sources.

Issued as interim guidance while formal University AI policy is finalized, these guidelines do not create new legal obligations. The guidelines operationalize existing requirements under UTEP, the UT System, and applicable federal and state frameworks, and apply them to the AI research context. The guidelines provide explicit guidance on five areas: mandatory disclosure of AI use in research outputs; data governance requirements that restrict sensitive, controlled, or regulated data to approved and contractually protected environments; human accountability for all AI-assisted outputs; compliance with IRB/HRPP, export control, sponsor, and Data Use Agreement requirements as they apply to AI; and practical guidance on records retention, intellectual property, and research security. The guidelines also provide a clear institutional risk-tolerance statement, distinguishing uses that are “strongly discouraged” from those that are “not permitted” and include unified disclosure language sufficient for use in scholarly publications as well as grant proposals.

Critically, these guidelines regulate risk, compliance, and disclosure; they do not prescribe research questions or methodologies, and all AI-use decisions remain scholarly decisions, subject to the existing norms of research integrity and the professional judgment of the researcher.

2. Purpose

These guidelines provide University-wide guidance for the responsible, secure, and ethical use of Artificial Intelligence (AI) and Generative AI (GenAI) in research conducted at The University of Texas at El Paso (UTEP). They are intended to support researchers in the design, conduct, analysis, and dissemination of research when AI systems or AI-enabled tools are used as part of the research process.

The following guidelines serve as interim guidance while policy for the ethical use of AI is being drafted. They do not create new legal obligations independent of existing University, UT System, or regulatory requirements. Rather, they are designed to help the UTEP research community understand and apply those existing obligations in the context of AI-enabled research, and to promote good research practice consistent with the UTEP Handbook of Operating Procedures (HOP), UTEP Information Security Standards, UT System policies, and applicable state and federal laws.

3. Scope

These guidelines are intended for all individuals engaged in research activities conducted under the auspices of UTEP, regardless of funding source(s). This includes faculty, staff, students, postdoctoral researchers, visiting scholars, and contractors. Graduate students

conducting thesis or dissertation research are fully covered; see Section 17 for thesis and dissertation-specific guidance and examples.

The guidelines apply in all cases where AI systems or AI-enabled tools are used as part of the research process, including their development, acquisition, deployment, evaluation, or integration into research workflows. They also apply to AI use within research-related administrative activities that directly support the conduct, oversight, or management of research.

These guidelines are not intended to apply to:

- Routine administrative or office productivity uses of AI that are not directly related to the conduct of research.
- Instructional or classroom uses of AI that are not part of a research activity.
- Personal or non-research uses of AI by UTEP personnel conducted outside the scope of university research activities.

4. Relationship to Existing Policies and Governing Authorities

These guidelines are issued under the authority of the UTEP Office of the Vice President for Research (OVPB) to support responsible and compliant use of AI in research. They are intended to complement — not replace or supersede — existing UT System policies, UTEP institutional policies, and applicable state and federal laws. In all cases, governing law, regulation, UT System policy, and UTEP policy take precedence over these guidelines.

Where these guidelines use the language of obligation (e.g., "researchers should," "it is recommended that"), this reflects either: (a) a best practice strongly encouraged by UTEP, or (b) an obligation that already exists under one of the governing authorities listed below and that these guidelines are helping to operationalize. Researchers are encouraged to consult the relevant governing authority directly for binding requirements.

These guidelines are informed by, and should be interpreted consistently with, the following authorities:

- UT System Policy UTS 165, Information Resources Use and Security, including UTEP Information Security Office (ISO) Standards (Standard 2: Acceptable Use; Standard 9: Data Classification; Standard 11: Safeguarding Data; Standard 22: Vendor and Third Party Controls; and Standard 26: Acceptable Use of AI Tools);
- UT System Policy UTS 173, Export Controls;
- UT System Policy UTS 115, Records and Information Management;
- UTEP Handbook of Operating Procedures (HOP), including the Research Integrity Policy;
- UTEP Institutional Review Board (IRB) / Human Research Protection Program (HRPP) requirements, where applicable;
- UT System HIPAA privacy policies governing Protected Health Information (PHI), where applicable; and
- Applicable local, state, and federal laws and regulations; and

- Texas Administrative Code Title 1, Part 10, Chapter 219 (TAC 219), Code of Ethics and Ethical Principles of Artificial Intelligence, which establishes state-level AI governance requirements applicable to Texas state agencies and their contractors, including requirements for an AI Risk Officer, heightened scrutiny AI system controls, human-in-the-loop oversight, public disclosure using DIR templates, and annual AI training with an initial compliance deadline of August 31, 2026.

5. Relationship to Academic Freedom

These guidelines regulate risk, compliance, and disclosure and are not intended to regulate intellectual direction. The selection of research questions, methodologies, and analytical approaches remains entirely within the discretion of UTEP researchers, consistent with the University's longstanding commitment to academic freedom. The guidelines address data security, tool approval, disclosure obligations, sponsor compliance, and records retention, which are all operational and administrative matters, and not scholarly ones. All AI-use decisions are, at their core, scholarly decisions, subject to the professional judgment of the researcher and to the existing standards of research integrity and disciplinary practice that govern the conduct of research at UTEP. Nothing in these guidelines should be read to direct research outcomes or to constrain the independent pursuit of knowledge.

6. Definitions

For purposes of these guidelines, the following terms are used as described below. These definitions are intended to support shared understanding and are not intended to replace or supersede definitions established in applicable laws, regulations, or University or UT System policies.

AI System: An engineered or machine-based system designed to generate outputs — such as text, images, predictions, recommendations, or decisions — using computational techniques that may operate with varying levels of autonomy.

Heightened Scrutiny AI System (HSAI): An AI system that makes consequential decisions having a material legal or significant effect on access to government services or benefits. Under TAC 219, HSAI systems are subject to additional controls and oversight requirements beyond those applicable to standard AI systems, including enhanced risk assessments, mandatory human-in-the-loop review, and specific documentation and disclosure obligations. Researchers and administrators deploying AI systems that may qualify as HSAI under TAC 219 should consult the UTEP ISO and the AI Risk Officer before deployment.

Generative AI (GenAI): A category of AI systems or models designed to generate new content based on training data and user inputs or prompts, including but not limited to text, images, audio, video, or code.

Large Language Model (LLM): A class of generative AI models that use machine learning techniques to process, understand, and generate human language. These models are typically trained on large corpora of text and may be used to recognize, summarize, translate, predict, and generate language-based content.

AI-Enabled Tool: A software application, platform, or service that incorporates AI capabilities to assist with tasks such as content generation, data analysis, pattern

recognition, decision support, or automation. AI-enabled tools may be standalone systems or embedded within broader research, computing, or administrative environments.

Note: AI capabilities are increasingly embedded in everyday tools, including word processors, coding environments, search engines, email clients, and scientific instruments, often without explicit user invocation. The obligations in these guidelines apply to deliberate, material uses of AI (see “Substantive AI Use” below), not to ambient or incidental AI features that operate in the background of standard software environments.

Substantive AI Use: A deliberate engagement with an AI system or AI-enabled tool that meaningfully shapes a research output, analytical result, or research-related work product. This use includes, but is not limited to: querying a large language model to generate, summarize, or revise text that is incorporated into a manuscript, proposal, or report; using an AI tool to perform data analysis, pattern recognition, or predictive modeling that contributes to research findings; using AI to generate code that forms part of a research software product; and using AI systems to support hypothesis generation or experimental design. Substantive AI use does not include: AI-assisted autocomplete or inline code suggestions in standard development environments where the researcher reviews and accepts or rejects each suggestion; AI-powered grammar or style suggestions in word processing software; AI-enhanced search results in standard web or database search tools; or AI features embedded in scientific instrumentation that operate as part of the instrument's standard data processing pipeline and are disclosed through standard methods reporting. Researchers who are uncertain whether a particular use qualifies as substantive should treat it as substantive and apply the relevant disclosure and records obligations.

Research Records: Documentation or materials that embody the data, methods, or analyses of research activities. Research records may include, but are not limited to, proposals, laboratory records, datasets, software code and scripts, figures, manuscripts, theses/dissertations, presentations, and other research outputs.

Restricted Data: Information or data used in university research that is subject to limitations on its access, use, or disclosure based on contractual obligations, sponsor requirements, applicable law or regulations, or UTEP or UT System policy.

Sensitive Data: Information relating to an identifiable individual that identifies, or can reasonably be used to identify, an individual — directly or indirectly, including in combination with other data. This may include identifiers such as names, identification numbers, location data, online identifiers, or other attributes associated with an individual's identity.

Human Subjects Data: Data obtained through interaction or intervention with living individuals, or identifiable private information about living individuals, that is collected, generated, analyzed, or used as part of research involving human subjects. Human subjects data may include direct identifiers, coded data, or data that could reasonably be used to re-identify individuals, consistent with applicable IRB and HRPP requirements.

Research-Related Administrative Activity: Administrative or operational activities that directly support the planning, conduct, oversight, management, or reporting of research. Examples include research administration, compliance activities, data management, proposal preparation, project management, sponsor reporting, and coordination of research personnel or resources.

This term is used to clarify that these guidelines extend to AI use in administrative support of research, such as using AI to draft a progress report or manage research data, and not solely to AI used in the scientific or scholarly work itself. Routine office or institutional administrative activities that have no direct connection to a specific research project are outside this definition and outside the scope of these guidelines.

7. Core Principles

The following principles are intended to guide researchers in making sound, ethical decisions about AI use in their work. They reflect broadly accepted norms in research integrity and responsible AI practice at peer institutions.

1. Attribution & Transparency — The use of AI in proposals, protocols, manuscripts, software, datasets, and presentations shall be clearly disclosed in citations, references, footnotes, or appendices. AI tools are not authors or inventors, and their contributions should be distinguished from human intellectual work.

2. Human Accountability — Researchers who use AI in their work remain solely responsible for the validity, accuracy, and ethical integrity of any output they share, submit for publication, present, or otherwise disseminate. AI outputs should always be reviewed and validated by a qualified human researcher before use. Note: TAC 219 independently requires a “human in the loop” for AI decision-making processes subject to its scope. Researchers and administrators deploying AI systems covered by TAC 219 should ensure that meaningful human oversight is built into the workflow, not merely applied after the fact.

3. Privacy, Security & Non-Discrimination — Researchers shall handle data in accordance with UTEP and agency data classifications and required safeguards; assess and mitigate bias consistent with the risk level and disciplinary standards; and remain attentive to privacy considerations as they pertain to themselves, their collaborators, and their study participants.

4. Compliance-by-Design — Researchers are encouraged to select and configure AI tools in a manner consistent with UTEP ISO controls and UT System policies from the outset of a project, rather than addressing compliance concerns after the fact. It is prudent to assume that inputs to third-party AI services may be retained or exposed unless contractually protected.

5. Inclusivity & Equity — The development and deployment of AI should prioritize inclusivity and accessibility, address potential biases, and seek equitable outcomes for all communities. AI should not be used in ways that exacerbate biases, exploit communities, or undermine human dignity. Note: TAC 219 requires that when members of the public interact with a covered AI system, disclosure must be made using the DIR-prescribed template, and a mechanism for redress must be available to any individual who believes they have been discriminated against by an AI system. These specific requirements are addressed in Section 13 (Suggested Unified Disclosure Language). Researchers and administrators whose AI deployments involve public-facing systems should consult the AI Risk Officer and the UTEP ISO to ensure compliance.

7.1. Institutional Risk-Tolerance Statement

AI tools offer real and growing benefits to researchers, accelerating literature synthesis, enabling new forms of data analysis, and supporting the kind of iterative, exploratory work that allows for impactful scholarship. UTEP encourages the responsible use of AI tools. At

the same time, the potential for adverse consequences to data subjects, research integrity, the institution's legal/ contractual obligations, and potentially to national security, requires safeguards adequate to address the risk involved. These guidelines are designed to support responsible AI use while managing the risks that come with it.

For the purpose of interpreting these guidelines, three distinctions are relevant. Uses described as **"permissible"** are consistent with UTEP policy when applicable disclosure, validation, and data-handling requirements are met. Permissible uses represent the normal operating space for AI-assisted research at UTEP. Uses described as **"strongly discouraged"** carry meaningful institutional, legal, or reputational risk; researchers are expected to avoid "strongly discouraged" uses unless there is compelling justification and prior consultation with UTEP AI in Research Committee (AIRC-UTEP) or the relevant compliance office. Uses described as **"not permitted"** are prohibited outright, as they are inconsistent with governing law, sponsor requirements, or UTEP/UT System policy. Engaging in "not permitted" uses may subject the researcher to consequences under those existing frameworks. Chairs and deans interpreting these distinctions should treat "not permitted" as a hard boundary and "strongly discouraged" as a risk-management signal that warrants case-by-case assessment before proceeding.

Researchers uncertain about which tier applies to their situation — particularly those working in data-heavy or computational contexts where the boundaries are not always obvious — should work through the following questions before proceeding: (1) Does the data involved qualify as regulated, controlled, or sensitive under UTEP Standard 9 or applicable law? (2) Is the data governed by a Data Use Agreement, sponsor agreement, or IRB protocol? (3) Has the proposed AI tool been reviewed and approved by the UTEP ISO? (4) Does the tool have a Data Protection Addendum or, where PHI is involved, a Business Associate Agreement? A "yes," "unclear," or "unknown" answer to any of these questions means the use should be treated as at minimum "strongly discouraged," and the researcher must consult AIRC-UTEP or the relevant compliance office before proceeding.

8. Permissible and Discouraged Uses

Conversely, using an LLM to generate a draft manuscript section, using AI to perform statistical analysis contributing to research findings, or using an AI tool to generate code that forms a material part of a research software product are all considered substantive AI Uses subject to the full obligations of these guidelines. Researchers uncertain whether a specific use is material should consult AIRC-UTEP or treat it as material and apply the relevant obligations as a conservative default. Note also that enterprise tools contracted by UTEP, such as Microsoft Copilot (note: enterprise data protections for tools such as Microsoft Copilot apply only when accessed through university credentials, i.e., @utep.edu accounts; additionally, sensitive or confidential research data should not be entered into the prompt of any AI system when web search or internet retrieval features are enabled, as data submitted in that context may be transmitted to external search providers and is not protected by the university's enterprise agreements)

The following uses of AI in research are generally permissible, subject to applicable disclosure, validation, and data-handling requirements established in these guidelines and in governing policies:

- Helping to generate research ideas, hypotheses, and approaches;
- Supporting research activities, including data analysis, pattern recognition, figure drafting, and predictive modeling;

- Assisting with published literature review and triage;
- Coding assistance, debugging, and software development support; and
- Drafting, editing, and summarizing research-related text, provided all AI contributions are disclosed and human-reviewed. This applies to the researcher's own intellectual contributions; using a collaborator's text or intellectual work with an AI tool requires that collaborator's explicit prior consent.

The following uses are strongly discouraged and, in most cases, inconsistent with research integrity standards, governing institutional policies, or applicable law. Researchers engaging in these activities may be subject to consequences under those existing governing frameworks:

- Uploading confidential or controlled data — including Protected Health Information (PHI), student records protected under the Family Educational Rights and Privacy Act (FERPA), Controlled Unclassified Information (CUI), export-controlled data, or proprietary sponsor data — into non-approved AI tools. This is inconsistent with UTEP ISO Standards and, depending on the data type, may violate federal law or sponsored research agreements,
- Using AI to fabricate or falsify research data, results, processes, or outputs. This constitutes research misconduct under the UTEP HOP Research Integrity Policy;
- Submitting AI-generated text, citations, or analysis as original work without disclosure or human validation; and
- Processing confidential peer-review materials through AI unless the venue or journal explicitly permits it and the tool has been reviewed through the AI Tool Review process. For clarity, “peer-review materials” here refers to manuscripts, proposals, or other submissions provided to a researcher in their capacity as a referee or reviewer for a journal, conference, or funding agency — not to published literature review, which is addressed separately as a permissible use above.

Scenarios & Examples

The examples below illustrate clear cases at each end of the spectrum. Researchers working in computational, data-intensive, or interdisciplinary contexts will often encounter situations that do not map cleanly onto a single tier. In those cases, the diagnostic questions in Section 7.1 provide a starting point, but there is no substitute for direct consultation with AIRC-UTEP or the relevant compliance office before proceeding. The scenarios here are illustrative, not exhaustive.

✓ Permissible: AI-Assisted Pattern Recognition on Public Data: A computational geoscience researcher uses a UTEP-approved AI tool to assist with pattern recognition in a large, publicly available seismic dataset. No controlled, sensitive, or DUA-governed data is involved. The tool is on the UTEP ISO Approved AI Tools list. The researcher validates outputs manually, discloses AI use in the methods section, and retains prompt logs as part of the research record. All four diagnostic questions in Section 7.1 are answered “no,” confirming this falls squarely in the permissible tier.

⚠ Strongly Discouraged: Consumer AI Tool with DUA-Governed Data: A faculty member wishes to use a consumer-grade AI platform to analyze a dataset governed by a sponsor Data Use Agreement (DUA) that does not address AI use. The DUA's lack of mention of AI use creates meaningful contractual and legal ambiguity — this is not a hard prohibition, but it is not a green light either. The researcher must consult the UTEP Office of Sponsored Projects before proceeding. If the ambiguity cannot be

resolved internally, the researcher must seek written clarification from the data provider before any AI tool is used. Proceeding without that consultation is strongly discouraged.

✗ Not Permitted: PHI Uploaded to a Non-Approved Consumer AI Tool: A researcher uploads clinical interview data containing Protected Health Information (PHI) to a free, consumer-grade AI chatbot to assist with analysis. The tool has no Business Associate Agreement, no ISO approval, and is not covered by the IRB protocol. This is a hard stop: regardless of the researcher's intent, this use is inconsistent with HIPAA, UTEP ISO Standards, and the IRB-approved protocol, and may expose the researcher and institution to federal penalties and legal liability. No consultation or justification can move this use out of the "not permitted" tier.

✓ Valid Scenario — Literature Review Support: A graduate student uses a UTEP-approved GenAI tool to summarize recent publications related to cybersecurity threat detection. They verify every reference, ensure accuracy, and disclose AI use in the methods section of their thesis.

✓ Valid Scenario — Code Optimization: A research scientist uploads non-sensitive, synthetic example code into an approved AI coding assistant to debug performance issues in a simulation script — without sharing research data or proprietary sponsor code.

✗ Invalid Scenario — AI-Generated Citations and Fabricated Content: A PI copies an AI-generated manuscript section (including invented citations) into a paper without verification. This violates research integrity standards and the human-accountability principle, and may constitute research misconduct under the UTEP HOP.

✗ Invalid Scenario — Undisclosed AI Drafting: A postdoctoral researcher uses a GenAI tool to write the results and discussion sections of a journal submission without disclosing this to co-authors or in the manuscript. This is inconsistent with authorship standards and disclosure obligations under these guidelines and applicable journal policies.

9. Data Governance & Classification

All University data used with AI should be classified and handled in accordance with UTEP Standard 9 (Data Classification) and related ISO standards. Confidential or controlled data — including PHI, FERPA student records, CUI, export-controlled technical data, and unpublished sponsor-proprietary data — should only be processed in UTEP-approved, contractually protected environments and tools. Researchers are strongly discouraged from disclosing such data to public or consumer AI services, as this is inconsistent with UTEP ISO Standards and may violate applicable law or sponsored-research agreements.

Before using any AI tool with data classified above the "public" tier, researchers are encouraged to confirm the tool's approval status through the UTEP ISO Approved AI Tools list or to submit for review through the AI Tool Review intake process. When uncertain about data classification, researchers should consult the UTEP ISO or their department's research compliance officer.

Researchers engaged in collaborative or multi-institutional projects must ensure that AI use across all participating entities comply with applicable institutional policies, sponsor requirements, and data governance obligations. Tool approval is institution-specific: a tool

reviewed and approved by the UTEP ISO is not automatically cleared for use by collaborating personnel at a partner institution, nor is a partner institution's approved tool automatically acceptable for use with UTEP data. Where collaborators from different institutions will share data or jointly use AI tools, these differences should be addressed contractually through the collaboration agreement, Data Use Agreement, or a dedicated AI use addendum before data is shared or analysis begins. Researchers should not assume that silence in a collaboration agreement on AI use constitutes permission. The UTEP Office of Sponsored Projects can assist in negotiating explicit AI use provisions — covering permitted tools, data handling requirements, and disclosure obligations at the agreement drafting stage, when it is far easier to resolve than after the collaboration is underway.

Scenarios & Examples

✓ **Valid Scenario — Anonymized Data Processing:** A research team de-identifies a human-behavior dataset and analyzes it in a UTEP-approved, contractually protected AI platform. The IRB protocol covers AI use and re-identification risk.

✓ **Valid Scenario — IRB-Approved AI Transcription:** A PI uses a PHI/FERPA-compliant, ISO-reviewed transcription AI tool to transcribe interview recordings under an approved IRB protocol. All files remain in a secure UTEP environment.

✗ **Invalid Scenario — Uploading Regulated Data to a Public AI Tool:** A researcher pastes lab notes containing PHI into a consumer AI chatbot to "clean up wording." This is inconsistent with HIPAA requirements, UTEP data-governance standards, and the terms of the IRB-approved protocol, and may expose the researcher and institution to legal liability.

9.1. Data Use Agreements (DUAs)

Many research datasets are governed by Data Use Agreements (DUAs) which are legally binding contracts between UTEP and a data provider. DUAs specify, in detail, how data may be accessed, used, stored, shared, and disposed of. With the introduction of AI, DUAs also specify whether and how AI tools may be used in connection with the covered data. DUAs control. Where a DUA is agreed upon between researchers and data providers, researchers are bound by the terms of the DUA. General permissions described elsewhere in these guidelines do not override a DUA's restrictions.

Prior to using any AI tool to process data governed by a DUA, researchers must:

1. review the DUA carefully for provisions addressing permissible use environments, data transfer restrictions, third-party access, approved analytical methods, security requirements, personnel authorizations, publication restrictions, and breach notification obligations, all of which are standard DUA elements that directly bear on AI use.
2. consult with the UTEP Office of Sponsored Projects or the Office of General Counsel before proceeding, in the case where the DUA is not clear on the use of AI tool or when the terms are ambiguous. It is important to consider that many DUAs were agreed upon before AI-assisted research became common and might not contemplate its use. Ambiguity must be resolved with the data provider and not assumed away.
3. obtain any required written approval or amendment from the data provider before any AI tool is used to access the covered data.

Security is a separate but equally important obligation. DUA-governed data must be processed only in environments that satisfy the security requirements specified in the agreement. Researchers should confirm with the UTEP ISO that any proposed AI environment meets those controls before use. For guidance on the tool approval process, see Section 11 (AI Tool Approval, Procurement & Vendor Risk).

Finally, for researchers who are currently negotiating new DUAs, use of AI with the data in question should take place at the DUA's drafting stage, not after the data has been provided. UTEP's Office of Sponsored Projects can assist in incorporating explicit AI use provisions, covering permitted tools, security requirements, and disclosure obligations as to not cause ambiguity to become an obstacle when analysis begins.

10. Human Subjects, Re-identification Risk, and HRPP Considerations

Projects that involve human participants or identifiable private information are required by federal regulation and UTEP IRB/HRPP policy to engage the IRB at the design stage. These guidelines encourage researchers to address the role of AI specifically within that required process. Protocols and informed consent materials should describe the role of AI, data flows, storage and retention practices, potential re-identification risks, potential group-level harms, and the nature of human oversight applied to AI outputs.

Informed consent documents should explain in plain language how AI will be used in the study, so that participants can make an informed decision about whether to participate. Researchers are encouraged to be aware of the risks posed by introducing even de-identified data into AI tools — including re-identification risk, confidentiality breaches, and the limitations of AI-generated analysis — and to apply secondary validation of AI outputs where appropriate. Where applicable, researchers should comply with controlled-access and data repository requirements (e.g., NIH Controlled Access Data Repository (CADR)).

Scenarios & Examples

✓ **Valid Scenario — AI Sentiment Analysis Under IRB Oversight:** A psychology research team uses an approved AI sentiment-analysis tool to process de-identified text responses from a clinical study. The IRB protocol explicitly covers AI-assisted analysis, describes re-identification risks and mitigations, and specifies that all data remain in a UTEP-managed secure environment. The informed consent form explains in lay terms that computerized analysis tools will be used. The PI manually validates a random sample of AI outputs.

✓ **Valid Scenario — Anonymized Survey Analysis with Proper Approvals:** A public-health researcher submits the AI tool through the AI Tool Review intake before beginning analysis of survey responses containing indirect identifiers. An ISO evaluation confirms U.S. data residency, the vendor contract prohibits training on user inputs, and SSO/MFA enforcement is confirmed. The IRB protocol is amended to describe the AI's analytical role before analysis begins.

✗ **Invalid Scenario — AI Analysis Added After IRB Approval Without Amendment:** A PI receives IRB approval for a study involving identifiable interview data, then uses a third-party AI tool to automate transcript analysis without submitting an IRB amendment or verifying the tool meets UTEP data-security requirements. This is inconsistent with HRPP requirements, UTEP ISO data-governance standards, and the informed-consent agreement made with participants.

11. AI Tool Approval, Procurement & Vendor Risk

Researchers are strongly encouraged to seek formal pre-use review through the UTEP AI Tool Review (the AI Tool Review intake process is administered by the UTEP ISO; documentation of the review criteria, submission process, and approved tools list is available through the UTEP ISO at <https://www.utep.edu/information-resources/iso/>; researchers should confirm the current intake URL directly with the ISO, as this resource may be updated) intake process for any AI tool that will process confidential or controlled data, or that could affect individual rights or safety. Where the data involved is subject to specific legal or contractual protections (e.g., PHI under HIPAA, CUI under federal law, or data governed by a sponsored-research agreement), such review is required under those governing frameworks and these guidelines help researchers fulfill that obligation.

The review process evaluates: security controls under UTEP ISO Standard 22; data residency; training behavior (whether the vendor uses submitted prompts or data to train or improve its models); logging and telemetry practices; access controls including Single Sign-On (SSO) and Multi-Factor Authentication (MFA); encryption standards; and required contractual protections. Contractual protections should include, at minimum, a Data Protection Addendum (DPA) governing how the vendor handles university data, and — where the tool will process Protected Health Information (PHI) — a Business Associate Agreement (BAA) as required under HIPAA. Only tools that have been reviewed and approved or formally contracted by the ISO should be used for regulated or unpublished research data.

Scenarios & Examples

✓ **Valid Scenario — Approved Tool for Sensitive Survey Data:** A public-health researcher submits an AI tool through the AI Tool Review intake and receives ISO approval confirming U.S. data residency and a DPA prohibiting use of inputs for model training. SSO/MFA is confirmed. The IRB protocol is amended to describe the AI's role. Analysis proceeds only after written approval is received.

✓ **Valid Scenario — Use of an ISO-Listed Coding Assistant:** A computer-science PhD student uses a coding assistant on UTEP's Approved AI Tools list to develop simulation software for an NSF project. No sponsor-proprietary data or export-controlled algorithms are entered. The student discloses AI assistance in the publication's methods section and retains prompt logs as part of the research record.

✗ **Invalid Scenario — Unapproved Tool with Sponsor-Proprietary Data:** A research team uses a consumer-grade AI platform — without ISO review or a DPA — to process unpublished sponsor-proprietary technical specifications. This is inconsistent with the AI Tool Approval guidance in these guidelines, may violate the terms of the sponsored-research agreement, and could expose the university to legal liability.

12. Export Controls & Controlled Unclassified Information (CUI)

Federal law — including the Export Administration Regulations (EAR) and the International Traffic in Arms Regulations (ITAR) — prohibits the transfer of export-controlled technical data to unauthorized recipients or environments. Similarly, CUI handling requirements under federal policy restrict how controlled information may be processed and stored. These legal obligations apply regardless of whether AI tools are involved; these guidelines simply highlight how they apply in the AI context.

Researchers working with export-controlled data or CUI should use only university-authorized (or funding agency provided) computing environments that meet applicable federal requirements and should never enter such data into AI tools that have not been specifically authorized for that purpose. Before using any AI tool in connection with export-controlled data or CUI, researchers shall contact the UTEP Export Control Officer for guidance. Where required, a Technology Control Plan (TCP) — a formal document describing how controlled data or technology will be accessed, used, and safeguarded — should be in place prior to any AI-assisted activity involving such information.

Scenarios & Examples

✓ **Valid Scenario — CUI Research with an Authorized AI Environment:** An engineering faculty member working on a DoD-funded project involving CUI contacts the UTEP Export Control Officer before selecting an AI tool. After review, the team is directed to a government-provided list of requirements for computing environments. The Export Control Office assists in developing a Technology Control Plan. No CUI is entered into any commercial or cloud-based AI service.

✗ **Invalid Scenario — Export-Controlled Data Entered into a Public AI Tool:** A postdoctoral researcher working with export-controlled technical data pastes it into a commercial AI chatbot to help draft a technical report, without consulting the Export Control Officer or verifying whether the tool is authorized. This may constitute a violation of EAR/ITAR, exposing the university and individual to federal penalties, and is inconsistent with UTEP policy and UTS 173.

13. Suggested Unified Disclosure Language

Disclosure of AI use in research proposals, manuscripts, and other research outputs is required under these guidelines. However, these guidelines do not mandate a specific form of words. To support consistency across the institution and to reduce the burden on researchers who would otherwise draft disclosure language independently, UTEP recommends the unified disclosure statements below. Researchers should select and adapt the language most appropriate to their context. As always, publications or sponsor-specific disclosure requirements take precedence where more specific guidance is provided.

Sample Disclosure Language for Publications:

“The authors used [name of AI tool(s)] to assist with [specific task(s), e.g., literature search, data analysis, manuscript drafting, code development]. All AI-generated outputs were reviewed, validated, and revised by the authors, who take full responsibility for the accuracy and integrity of the work.”

Optional add-on for contexts where the venue or sponsor specifically requires a research-integrity affirmation: “AI tools were not used to fabricate data, generate fictitious citations, or replace human intellectual contribution.” Researchers should include this sentence only where required or requested by the venue; it is not part of the recommended default language, as most publication venues do not expect or require it.

Sample Disclosure Language For grant proposals:

“Artificial intelligence (AI) tools were used in the preparation of this proposal for the following purpose(s): [describe use, e.g., editing for clarity, literature synthesis, data visualization]. All substantive intellectual content, research design, and certifications

are the responsibility of the Principal Investigator(s). Use of these AI tools complies with applicable solicitation requirements and UTEP institutional policy.”

Note: Some federal solicitations prohibit or restrict AI use in proposals. Researchers must consult solicitation-specific instructions before using AI in proposal preparation. See Section 14 (Compliance and Sponsor-Facing Guidance) for agency-specific guidance.

Venue-Specific Disclosure Requirements

The unified disclosure language in Section 13 represents UTEP’s recommended baseline, which is the minimum requirement. Journals, conferences, and funding agencies increasingly maintain their own AI disclosure requirements that may be more specific, more stringent, or structurally different from UTEP’s recommended language. Where a venue’s requirements exist, they must be followed precisely. Researchers should consult the submission guidelines of every target venue before finalizing any disclosure statement and should not assume that UTEP’s unified language satisfies venue-specific requirements.

Venue requirements vary considerably. Some journals and conferences require disclosure only in an acknowledgments section; others require section-level citations identifying each portion of the paper where AI-generated content appears, accompanied by a brief explanation of how the tool was used. Some venues distinguish between AI used for substantive content generation and AI used for grammar or editing assistance, treating the latter as generally outside their disclosure policy. Double-anonymous submission requirements add a further constraint: disclosure language must not reveal author identity or funding sources. Researchers serving as referees or reviewers should also be aware that many venues now explicitly prohibit or restrict the use of AI tools in the peer-review process; where such restrictions exist, they are binding obligations, not recommendations.

14. Compliance and Sponsor-Facing Guidance

AI use in the context of sponsored research can be present at multiple stages, each with its own compliance obligations. These guidelines address three main uses of AI in sponsored research: 1) idea and hypothesis development - AI is used to generate or refine research questions, explore conceptual frameworks, or identify gaps in the literature; 2) proposal preparation - AI is used to assist with drafting, editing, literature review, or budget justification; and 3) research conduct - AI is integrated into data collection, analysis, modeling, or reporting. Sponsor requirements, certification obligations, disclosure expectations, and intellectual property considerations differ across these stages, and researchers should treat them accordingly. Sponsor requirements, disclosure expectations, and intellectual property considerations differ across these three stages, and researchers should treat them accordingly. In addition to guidance on each stage, this section includes a dedicated subsection (14.3) on PI certification responsibility, which is a cross-cutting obligation that applies regardless of which stage AI is used.

14.1. AI Use in Idea and Hypothesis Development

Using AI to generate candidate research questions, explore theoretical frameworks, or identify promising directions in the literature is permissible, and in some case encouraged. However, this requires particular attention to intellectual property. As discussed in Section 21, federal law conditions patentability and copyrightability on human intellectual contribution, and AI participation in the conceptualization of a research idea can affect the protectability of resulting inventions or works. Researchers who use AI during ideation are strongly encouraged to document that use from the outset, recording what was prompted,

what the AI produced, and how human judgment was applied in developing the idea further. Researchers with questions about how AI-assisted ideation affects their IP position should consult the UTEP Office of Technology Commercialization before the research progresses.

14.2. AI Use in Proposal Development - Agency-Specific Positions

Federal sponsors are actively evolving their positions on AI use in proposal preparation, and those positions are not uniform. Researchers must consult the specific solicitation and any applicable policy notices before using AI in proposal development. General awareness of an agency's position is not a substitute for carefully reading the solicitation, particularly as agency guidance in this area is rapidly changing. Below is a description of the current positions of some of the major funding agencies.

The National Science Foundation (NSF) has taken the most permissive position among the major research funders. Under PAPPG 24-1, AI use in proposal preparation is permitted, provided the extent and manner of that use are disclosed within the relevant proposal section. Failure to disclose is treated as misrepresentation under the certifications signed at submission. The PI remains responsible for the accuracy and integrity of all proposal content, including content developed with AI assistance.

The National Institutes of Health (NIH) has taken the most restrictive position. Effective September 25, 2025, NOT-OD-25-132 prohibits proposals that are substantially developed by AI, treating them as non-original and ineligible for funding. NIH uses AI-detection tools, and post-award detection of AI-generated content can result in misconduct referrals, cost disallowances, grant suspension, or termination. Limited AI use in specific aspects of proposal preparation may be acceptable, but researchers should approach NIH proposals with significant caution and verify current policy notices before submitting.

The Department of Defense (DoD) and the Department of Energy (DOE) have not issued specific guidance on AI use in grant proposal preparation as of the date of these guidelines. DoD's recent AI policy activity, including the January 2026 AI Acceleration Strategy, has been directed at internal military AI deployment rather than at how external researchers use AI when preparing grant proposals. DOE's AI policy activity has likewise focused on internal workforce deployment and its Genesis Mission science initiative rather than on regulating how applicants use AI in proposals to DOE program offices and national laboratories. For both agencies, requirements vary by program and contracting office, and individual solicitations may address AI use explicitly or remain silent. Researchers working on DoD- or DOE-funded programs should read each solicitation carefully and consult the UTEP Office of Sponsored Projects before using AI in proposal preparation.

14.3. PI Certification Responsibility

Regardless of the extent to which AI tools are used in proposal preparation, proposal certification remains the sole responsibility of the Principal Investigator. By submitting a proposal, the PI certifies the accuracy, completeness, and integrity of all proposal content, including content that was drafted or edited with AI assistance. The use of AI does not transfer, reduce, or otherwise affect this certification responsibility. PIs who use AI in proposal preparation are encouraged to retain records of how AI was used, consistent with Section 15 (Records & Retention), and to ensure that all AI-generated content has been reviewed, validated, and accepted as their own scholarly work prior to submission.

14.4. AI Use During Research Conduct

Once an award is secured, the terms of the award take precedence over general institutional guidance. Award terms, data management plans, reporting requirements, and any technology control provisions govern how AI may be used in data collection, analysis, modeling, and reporting. Researchers should review award terms carefully before deploying AI tools in research activities associated with the grant. Where award terms are silent on AI, researchers should consult the UTEP Office of Sponsored Projects to determine whether sponsor clarification is needed. Disclosure of AI use in research outputs, including progress reports, final reports, and publications, is required under these guidelines and, increasingly, under the terms of the award itself. Records of AI use during research conduct should be maintained as part of the official project record, consistent with Section 15.

15. Records & Retention

When AI tools materially contribute to the research process or its outputs, researchers are encouraged to preserve sufficient records to allow the work to be understood, verified, and reproduced. This practice supports research integrity and is consistent with the record-keeping obligations that already exist under UT System Policy UTS 115 and applicable federal sponsor requirements. Records should include, at minimum: the AI tool name and version used; key prompts or configuration settings; a description of how the outputs were integrated into the research; and any human review or validation steps taken.

All records related to AI-assisted research should be treated as part of the official research record and managed, retained, and disposed of in accordance with UTS 115 and the UTEP Records Retention Schedule. Researchers should be aware that AI-related records stored within third-party platforms may be subject to automatic deletion and should take steps to preserve relevant records in university-controlled systems.

Reproducibility is an increasingly important dimension of records retention for AI-assisted research. Where AI tools contribute to research results, the records retained should be sufficient not only to document what was done, but to allow independent reproduction of those results. In computational and data-intensive fields, this expectation is already well-established: many conferences and journals now encourage or require the submission of reproducibility artifacts, such as code repositories, containerized environments, or step-by-step execution instructions, alongside the research paper. Researchers using AI tools to produce results should document the specific tool, version, configuration, and prompts used with sufficient precision that the analysis could be reproduced by a qualified third party. Where the AI tool itself is integral to reproducing the results, researchers should consider whether the tool or its configuration can be preserved or archived alongside other research artifacts. Researchers should be aware that AI tools and services can be discontinued or substantially changed by vendors, potentially making results irreproducible after the fact; where this risk is material, it should be noted in the methods section and addressed proactively in the research record.

Researchers should also be aware that AI-detection tools, increasingly used by journals, conferences, and funding agencies to screen submissions, produce false positives. Human-authored text can be flagged as AI-generated, particularly in technical writing styles. Detailed, UpToDate research records, including drafts, notes, version histories, and research journals, provide the most reliable defense against such findings. Researchers are encouraged to maintain a running record of their research process in sufficient details that, if a false positive arises, they can demonstrate the human origin of their work. This

is consistent with good research practice generally and reinforces the records retention obligations already described in this section.

Scenarios & Examples

✓ **Valid Scenario — Maintaining an AI Research Log:** A research team using AI-assisted data analysis maintains a dedicated log recording, for each use: the tool name and version, the date, a description of the task, key prompts or configuration settings, a summary of the AI output, and the human review steps taken. The log is stored in the project's official electronic research notebook alongside raw data, consistent with the UTEP Records Retention Schedule and UTS 115.

✗ **Invalid Scenario — No Documentation of AI-Generated Analysis:** A faculty member uses an AI tool to perform statistical modeling for a federally sponsored project but keeps no record of the model version, prompts, or outputs, and deletes the tool's conversation history after each session. When the sponsor requests a methods audit two years later, the team cannot reproduce or verify the AI-assisted analysis. This is inconsistent with UTS 115 record-keeping guidance and federal research-integrity standards, and may jeopardize the award.

16. Privacy & Protected Health Information (PHI)

HIPAA requires that PHI — any individually identifiable health data created, received, or used during research be processed only in environments covered by appropriate contractual protections. These guidelines encourage researchers to ensure that any AI platform used to process PHI is formally contracted by UTEP, covered by a Business Associate Agreement (BAA), and approved by the UTEP ISO. Where a research protocol involves PHI, the IRB should review and approve the protocol, including the role of any AI tools used to process that data, as required by applicable HRPP policy.

Researchers handling PHI are encouraged to apply the principle of minimum necessary access, limiting data access to only what is required to accomplish the specific research task and to maintain required audit logs. Any suspected unauthorized disclosure or breach should be reported promptly in accordance with UTEP privacy policies and applicable breach notification requirements, which derive from HIPAA and UT System policy.

Scenarios & Examples

✓ **Valid Scenario — AI-Assisted Clinical Note Analysis:** A clinical research team uses a UTEP-contracted, HIPAA-compliant AI platform — covered by a BAA and approved by the ISO, to assist with analysis of de-identified clinical notes. The IRB protocol covers the AI tool's role, minimum-necessary access controls are in place, and audit logs are maintained throughout the study.

✗ **Invalid Scenario — Consumer AI Tool Used for PHI Processing:** A researcher uses a free consumer AI chatbot to help summarize patient interview transcripts, reasoning that the data is "mostly de-identified." Even partially identifiable PHI processed in a non-approved tool without a BAA is inconsistent with HIPAA requirements, UTEP ISO standards, and the IRB-approved protocol.

17. Thesis and Dissertation Research

Graduate students conducting thesis or dissertation research are subject to all applicable provisions of these guidelines and to any governing institutional policies referenced herein.

AI use that materially contributes to data collection, analysis, or writing should be disclosed to the supervising faculty advisor, documented in the thesis or dissertation (typically in the methods section or a dedicated appendix), and addressed in the relevant IRB protocol where human-subjects data are involved.

Thesis and dissertation committees are encouraged to discuss AI use expectations with students at the outset of the research, and to ensure that any AI use is transparent and consistent with disciplinary standards and any applicable journal, degree-granting, or accreditation requirements. Students should consult the UTEP Graduate School guidelines and their committee chair before incorporating AI tools into their research workflows.

Scenarios & Examples

✓ **Valid Scenario — AI-Assisted Qualitative Coding in a Dissertation:** A doctoral student in sociology uses a UTEP-approved AI platform to assist with initial thematic coding of 80 interview transcripts. The student reviews, validates, and revises all AI-generated codes; documents the tool name, version, and prompts in a methodology appendix; and confirms the approach with the committee chair and IRB before beginning. The dissertation committee reviews and approves the disclosure, consistent with practices at peer institutions such as Michigan, Cornell, and UT Austin.

✓ **Valid Scenario — Grammar and Clarity Editing of a Thesis Draft:** A master's student uses an AI writing assistant to check grammar, clarity, and sentence structure in a draft thesis chapter. The student reviews every suggested change, retains full intellectual ownership of the arguments and findings, and acknowledges the tool in the thesis acknowledgments section. No research data or confidential content is entered into the tool.

✗ **Invalid Scenario — Undisclosed AI Authorship of Dissertation Chapters:** A doctoral student uses a generative AI tool to draft multiple dissertation chapters without disclosing this to the advisor or committee, and submits the work as original scholarship. This is inconsistent with academic integrity standards under the UTEP HOP Research Integrity Policy and may result in disciplinary action under those existing policies.

18. Foreign Influence, Conflicts of Interest, and Research Security

Obligations to disclose substantial foreign relationships — including significant financial interests, sponsored research agreements, or participation in foreign talent recruitment programs — arise from federal law, sponsor requirements, and UTEP policy. These guidelines encourage researchers to consider how AI tool selection and use may intersect with those existing disclosure obligations, and to seek guidance proactively.

AI tool procurement and use should be reviewed for risks arising from foreign ownership or control — and researchers should also verify that any proposed AI tool does not appear on the Texas Department of Information Resources (DIR) Prohibited and Covered Applications list, available at <https://dir.texas.gov/information-security/covered-applications-and-prohibited-technologies>, which identifies technology products prohibited for use by Texas state entities. Use of a DIR-prohibited tool is not permitted regardless of any other approval status. vendor ownership or control, particularly where tools will process sensitive, export-controlled, or federally funded research data. Researchers are

encouraged to contact the UTEP Export Control Officer for guidance on screening AI tools and collaborators for foreign influence concerns.

Scenarios & Examples

✓ **Valid Scenario — Proactive Vendor Risk Screening:** A research team evaluating an AI tool for use in a federally funded project contacts the UTEP ISO and Export Control Officer to screen the vendor for foreign ownership or control concerns before procurement. The vendor passes review and the tool proceeds through the standard AI Tool Review intake process.

✗ **Invalid Scenario — Undisclosed Use of a Foreign-Controlled AI Platform:** A PI uses a research AI platform owned and operated by a foreign entity for analysis of sensitive DoD-funded project data, without disclosing this to the Export Control Officer or seeking ISO review. This may be inconsistent with federal research security requirements, the terms of the sponsored-research agreement, and UTEP policy.

19. Training & Competency

Researchers using AI tools in the conduct of research are encouraged to ensure they have the knowledge and training appropriate to their role and the nature of the data involved. Several of the training types described below are already required under existing governing frameworks; these guidelines highlight how those requirements apply in the AI research context.

Responsible Conduct of Research (RCR): Principal Investigators (PIs) and key research personnel shall complete RCR training as required by their funding agency and UTEP policy, including any modules addressing the use of AI in research, data integrity, and attribution.

AI training

IRB / Human Research Protection Program (HRPP): Researchers whose projects involve human subjects shall complete applicable IRB and HRPP training before using AI tools that interact with, analyze, or process human subjects' data, consistent with IRB policy requirements.

Information Security Awareness: Researchers using AI tools are encouraged to complete UTEP ISO awareness training covering data classification, safeguarding requirements, and the acceptable use of AI tools, consistent with ISO Standard 26.

Export Control: Researchers working on projects involving export-controlled data, CUI, or foreign national collaborators shall complete UTEP export control training prior to using any AI tool in connection with that work, consistent with UTS 173.

Discipline-Specific Training: Researchers are encouraged to pursue AI literacy and ethics training specific to their field, as standards for AI use vary across disciplines and are evolving rapidly. Research units and colleges are encouraged to provide or facilitate access to such training on a regular basis.

Records of completed training may be requested by UTEP's AI in Research Committee (AIRC-UTEP) or relevant compliance offices as part of the AI Tool Review process or a sponsored project audit.

20. Governance & Questions

The UTEP AI in Research Committee (AIRC-UTEP), operating under the Office of the Vice President for Research (OVPR) in collaboration with the Information Security Office (ISO), Office of General Counsel (OGC), IRB/HRPP, Export Control Office, University Libraries, and academic Colleges, administers these guidelines. AIRC-UTEP maintains the Approved AI Tools list, publishes guidance on emerging AI use cases, and is available to assist researchers navigating complex or ambiguous situations.

Researchers with questions about how these guidelines apply to their work are encouraged to contact AIRC-UTEP or the relevant compliance office (ISO, IRB, Export Control, OGC) before proceeding. Where a researcher believes a departure from the guidance in these guidelines is warranted, they are encouraged to consult with AIRC-UTEP and document the rationale and any compensating measures taken.

20.1. Compliance and Consequences

Compliance with these guidelines depends primarily on researcher integrity, professional judgment, and the culture of transparency that UTEP is committed to fostering in its research community. These guidelines do not create a surveillance mechanism, and UTEP does not monitor AI use in real time. The disclosure, records retention, and tool approval obligations described throughout these guidelines are themselves the primary compliance mechanism: they create a documented trail that supports verification, audit, and reproducibility, and they reflect the professional and ethical commitments researchers make when they submit proposals, publish findings, report inventions, and engage in sponsored research.

These guidelines do not create new penalties or disciplinary mechanisms. Non-compliance that rises to the level of research misconduct, including fabrication, falsification, or undisclosed use of AI that misrepresents the nature of the research, is addressed under applicable federal regulations. Data security violations carry consequences under UTEP ISO Standards, UT System policy, and applicable state and federal law. Violations of sponsored-research agreement terms, including data governance and disclosure obligations, are governed by the terms of the award and may result in sponsor-imposed consequences including cost disallowances, award suspension, or termination. Export control violations are subject to federal penalties under EAR and ITAR. In each case, these guidelines help researchers understand and fulfill obligations that already exist under those governing frameworks; the frameworks themselves establish the consequences for non-compliance.

Looking ahead, AIRC-UTEP intends to work with the Office of Research Integrity, the Office of General Counsel, and relevant institutional stakeholders to incorporate language consistent with these guidelines into the UTEP Handbook of Operating Procedures (HoOP), specifically within the Research Integrity chapter. This will provide a durable, policy-level foundation for AI research governance as institutional AI policy matures and is formally adopted. In addition, UTEP will embed AI use disclosure affirmations within proposal submission workflows, making disclosure a routine, expected step in the submission process rather than a post-hoc obligation. These steps will strengthen the compliance infrastructure over time without creating new bureaucratic burdens on researchers acting in good faith.

21. Intellectual Property

Intellectual property resulting from research is protected in accordance with UT Rule 90101, based on analysis and applications managed by the UTEP Office of Technology Commercialization (OTC). Researchers are encouraged to document and report the use of AI in research to the OTC, as undisclosed AI contributions may affect the patentability or copyrightability of the resulting work under applicable U.S. law. Specifically, the use of AI during the discovery or ideation phase of a potential invention constitutes a disclosable event: researchers must report such use to the OTC at the time of invention disclosure, identifying what AI contributed and how human judgment shaped the final conception. Researchers should also be aware that submitting invention-related content to a third-party AI service may itself constitute a public disclosure under 35 U.S.C. § 102, starting the one-year statutory bar after which U.S. patent rights may be forfeited. Researchers are strongly encouraged to consult the OTC before using any AI tool in connection with unpublished inventive concepts.

Federal law — including 35 U.S.C. (patent) and 17 U.S.C. (copyright) — requires that humans form the basis of inventorship and authorship. Inventions that utilize AI should be annotated to distinguish what was generated by AI from what constitutes human intellectual contribution, particularly at the conception stage, consistent with guidance at 90 Fed. Reg. 54636 (Nov. 28, 2025). For copyright purposes (17 U.S.C. § 102(a)), human authorship is fundamental for protection; AI-generated content may be included in a work if it is clearly identified prior to submission. Software development projects that integrate AI-generated code require particular attention to these disclosure and annotation considerations. Researchers incorporating AI-generated code into research software should also be attentive to software license terms. AI coding tools generate output under varying license conditions, some tools assert no license restrictions on generated output, while others apply terms that may conflict with the researcher's intended use, sponsor requirements, open-source obligations, or UTEP's IP policies under UT Rule 90101. Researchers should review the terms of service and any applicable license terms of the AI coding tool before incorporating generated code into sponsored, publishable, or commercially relevant software, and should consult the UTEP Office of Technology Commercialization where license conflicts or uncertainty arise.

Scenarios & Examples

✓ **Valid Scenario — Disclosing AI Use in an Invention Disclosure:** A faculty member independently conceptualizes an experiment, conducts the research, and uses AI tools to assist with code development for results analysis. AI is also used for grammar editing during manuscript preparation. The researcher discloses the invention to the Office of Technology Commercialization and identifies the specific AI contributions. Under 90 Fed. Reg. 54636, this use is permissible because the human performed the conceptualization and experimental work, while AI served only as a tool for analysis and writing.

Note: The copyrightability of AI-assisted work, including software code, is a genuinely evolving area of law. Current U.S. Copyright Office guidance requires meaningful human authorship for protection, and purely AI-generated content without sufficient human creative contribution may not qualify. However, where a human researcher directs, selects, and meaningfully shapes AI-assisted output, the extent of copyright protection depends on the specific facts and the degree of human creative involvement. Researchers with questions about whether AI-assisted code or other

outputs are protectable should consult the UTEP Office of Technology Commercialization before filing or publishing.

X Invalid Scenario — AI Use in Conceptualization: A faculty member uses an AI tool during ideation to generate candidate molecular structures for a new drug compound, then independently evaluates, selects, and validates the relevant structures. Because ideation is central to the inventive process, the AI-generated conceptualization of targets could be characterized as machine thought rather than human conception, and therefore may not qualify as an invention of the human mind under 90 Fed. Reg. 54636 and 35 U.S.C. § 101. The OTC reviews the disclosure and determines patentability based on the extent of human-inventive contribution.

X Invalid Scenario — Failing to Disclose AI-Generated Software Components: A research team develops a software tool using AI-generated code that constitutes a significant portion of the codebase, and files for copyright registration without disclosing the AI-generated content. Under 17 U.S.C. § 102(a) and current U.S. Copyright Office guidance, AI-generated content without sufficient human authorship may not be protectable, a conclusion further supported by U.S. Copyright Office Circular 61. Failure to disclose is also inconsistent with UTEP's IP reporting obligations under UT Rule 90101 and may undermine the validity of the copyright claim.

22. Review Cycle

These guidelines will be reviewed at least annually and updated as needed to reflect changes in technology, applicable law, sponsor requirements, and institutional standards. The AIRC-UTEP is responsible for initiating and coordinating the review cycle. Feedback from the UTEP research community is welcomed and may be submitted to the VPR at any time.

23. Appendix A - References (Selected)

- 1) UTEP Information Resources Use & Security Policy and ISO Standards (incl. Standards 2, 9, 11, 22, 26): <https://www.utep.edu/information-resources/iso/policies/>
- 2) UT System UTS 165, Information Resources Use & Security: <https://www.utsystem.edu/sites/policy-library/policies/uts-165-information-resources-use-and-security-policy>
- 3) UTEP Human Subjects Research (IRB/HRPP): <https://www.utep.edu/research/human-subjects-research/>
- 4) UTEP IRB Information: <https://www.utep.edu/research/human-subjects-research/irb/>
- 5) UTEP HOP Research Integrity Policy: <https://www.utep.edu/hoop/section-4/research-integrity-policy.html>
- 6) UT System UTS 173, Export Controls: <https://advancing.utsystem.edu/sites/policy-library/policies/uts-173-export-controls>
- 7) UT System UTS 115, Records and Information Management: <https://www.utsystem.edu/sites/policy-library/policies/uts-115-records-and-information-management>
- 8) UT System HIPAA Privacy Policies: <https://www.utsystem.edu/offices/employee-benefits/hipaa-and-privacy>
- 9) UTEP Research & Innovation Policies (Export Control, Research Security, CUI, etc.): <https://www.utep.edu/research/office-of-sponsored-projects/policies/>
- 10) UTEP Data Governance (HOP Section X, Ch. 2): <https://www.utep.edu/hoop/section-10/ch-2-data-governance.html>
- 11) UTEP ISO — Protection of Sensitive Digital Research Data: <https://www.utep.edu/technologysupport/files/docs/utepprotectionofdigitalresearchdata.pdf>
- 12) UT Rule 90101, Intellectual Property: <https://www.utsystem.edu/board-of-regents/rules/90101-intellectual-property>
- 13) 35 U.S.C. § 101 (2018): <https://www.law.cornell.edu/uscode/text/35/101>
- 14) 17 U.S.C. § 101–103 (2012): <https://www.law.cornell.edu/uscode/text/17>
- 15) U.S. Copyright Office, Copyright and Artificial Intelligence, Part 2: Copyrightability (Jan. 29, 2025).
- 16) Thaler v. Perlmutter, No. 23-5233, 2025 WL 839178 (D.C. Cir. Mar. 18, 2025).
- 17) 90 Fed. Reg. 54636 (Nov. 28, 2025), USPTO Guidance on AI and Inventorship.
- 18) Notice to the Research Community on the Use of Generative AI in the NSF Merit Review Process (December 14, 2023), available at <https://www.nsf.gov/news/notice-to-the-research-community-on-ai>; formalized in PAPPG 24-1, Supplement 1 (2024).
- 19) NIH — NOT-OD-25-132: Supporting Fairness and Originality in NIH Research Applications (July 17, 2025), effective September 25, 2025, available at <https://grants.nih.gov/grants/guide/notice-files/NOT-OD-25-132.html>