AI Guidance Policy Checklist

Use this 61-item checklist to see where your institution's AI guidance policy (AIGP) stands. Check each item your policy covers, then share the gaps with your sponsored programs office, leadership, information security office, and Research Integrity Officer.

Important foundational decisions

Each item below directly impacts a policy's strictness, applicability, and interpretability. Consider each item carefully and deliberately carry your decision throughout the details and language you use in your policy.

- [ ] **Required or recommended:** decide whether people must follow the document (a formal policy or standard) or whether it only offers advice (principles or guidance). Advice is easier to update as AI changes; required rules can be enforced when broken.
- [ ] **Home of the rules:** a standalone AI policy, AI clauses added to existing policies, or a crosswalk that maps each existing policy to AI
- [ ] **Scope model:** one institution-wide policy, separate policies by function, or a system framework with campus policies
- [ ] **Tool model:** a single approved tool, a tool-by-data-risk matrix, or a "use vetted tools" principle
- [ ] **Disclosure model:** any material use, a percentage threshold, required wording, or encouraged only
- [ ] **Enforcement:** consequences for violations and whether AI use is monitored
- [ ] **Research integrity stance:** whether the misconduct policy names AI explicitly or relies on existing definitions

1. Governance and scope

- [ ] Name an owner office and a standing review group
- [ ] Define who is covered: faculty, staff, students, contractors, affiliates
- [ ] List which functions are covered, following the scope model: administration, teaching, research, or all
- [ ] Define AI terms, including generative and agentic AI
- [ ] Map existing policies that already apply: data classification, acceptable use, privacy, IP, misconduct
- [ ] Check federal, state, sponsor, and system mandates, such as executive orders, system directives, and institutional governance.

2. Tools and procurement

- [ ] Publish the approved tools that follow from the tool model, and where to find them
- [ ] State which data risk level each tool is approved for
- [ ] Set a request process for new tools, with a named form
- [ ] Require security and privacy review for AI vendors, such as HECVAT or SOC 2 Type II
- [ ] Require contract terms that bar training on institutional data and set retention and breach notice
- [ ] Cover AI features added to software already in use
- [ ] Decide on named bans
- [ ] Require institutional accounts for university work

3. Data protection

- [ ] Link AI rules to the existing data classification tiers
- [ ] State plainly which data may never go into AI tools, such as SSNs, PHI and FERPA protected data
- [ ] Set an approval path for exceptions, naming who signs off
- [ ] Require data minimization and synthetic data for testing
- [ ] Require data-owner consent before sharing non-public data
- [ ] Say how to report an accidental disclosure

4. Disclosure and attribution

- [ ] Say what a disclosure includes: tool, version, purpose, extent, date
- [ ] Say where the disclosure goes: acknowledgments, methods section, a footnote, or an internal note to a manager
- [ ] State that spelling and grammar help needs no disclosure, if so
- [ ] Point to citation style guidance

5. Human oversight and accountability

- [ ] State that users are responsible for all output they use
- [ ] Require human review for accuracy, citations and bias before use
- [ ] Bar AI-only decisions on grades, hiring, admissions, discipline and services
- [ ] Require documenting how AI outputs were used for sensitive work

6. Research integrity

- [ ] State that AI cannot be an author or co-investigator
- [ ] Name AI-assisted fabrication, falsification and plagiarism as misconduct, using the 42 CFR 93 and NSF definitions
- [ ] Include proposals, progress reports and records of AI use in the research record
- [ ] Require researchers to keep records of AI use: tool, version, prompts and unedited outputs
- [ ] Make the PI responsible for the team's AI use and for mentoring on it
- [ ] Add AI to responsible conduct of research training
- [ ] Involve the Research Integrity Officer, Chief Information Security Officer, or similar data security official in drafting the policy and in AI-related allegations if you have one
- [ ] Require documenting AI contributions to inventions and copyrighted works

7. Research and sponsored programs

- [ ] Clarify that the PI certifies AI-assisted proposal content; research administration staff cannot
- [ ] Prohibit AI in confidential peer review of proposals, in line with NIH and NSF notices
- [ ] Require checking sponsor and publisher AI rules
- [ ] Route human-subjects and clinical data through IRB
- [ ] Address controlled data: HIPAA, FERPA, Common Rule, export control
- [ ] Decide whether AI tool costs may be charged to awards
- [ ] Set a review step before research administration records enter an AI workflow

8. Special uses

- [ ] **Meeting AI:** require participant notice and consent, or ban unapproved note-takers
- [ ] **Agentic AI:** no credentials to institutional systems, sandbox testing, limits on actions
- [ ] **Intellectual property:** check tool terms for output ownership; no uploads of others' copyrighted work
- [ ] **Accessibility:** AI outputs must meet accessibility policy

9. Training and communication

- [ ] Offer AI literacy training, and decide whether it is required or annual
- [ ] Publish one central page for tools, rules and contacts
- [ ] Name a contact for questions and incidents

10. Maintenance

- [ ] Set a review date and cycle for the policy
- [ ] Label it a living document and track versions
- [ ] Collect community feedback before finalizing
- [ ] Re-review approved vendors on a schedule