

NATIONAL CONFERENCE OF ACADEMIC DEANS • PLENARY SESSION

AI Policy on a Real Campus

What We Learned Building Responsible Governance at Arkansas Tech

Ms. Angela Black & Dr. Sean Huss

Arkansas Tech University

Thursday, July 16, 2026

1:00 – 2:00 PM

Room 111/113, Brewer-Hegeman

OPENING

A Real Campus, Not a Finished Model

- No dedicated AI institute, large legal staff, in-house engineering team, or unlimited money
- Artificial intelligence moved faster than our normal policy process and entered through multiple doors at once
- Early solutions worked for a while, then created new problems of their own
- Leadership turnover exposed how fragile a temporary task force can be
- The early failures are not side notes to this case study: they are the case study

“Those are not side notes to the case study. Those are the case study.”

THE CENTRAL QUESTION

What Will AI Touch: What Must It Never Replace

AI IS LIKELY TO TOUCH

- Teaching, writing, and assessment
- Advising and research
- Hiring, budgeting, and procurement
- Marketing and information systems
- Cybersecurity and records
- Accessibility, spanning nearly every part of university life

AI MUST NEVER REPLACE

- Human responsibility
- Meaningful oversight
- Due process
- Professional judgment
- A clear path back to a person when a decision affects someone's education, employment, opportunities, rights, or access to services

ORIGINS

A Narrow Charge, Fall 2024

- Task force created in Fall 2024 by the Associate Vice President for Academic Affairs
- The charge was mostly instructional: investigate student use, prepare the university, draft syllabus language and policy
- Work organized into three areas: ethics and policy, teaching and education, research
- The three-part structure contained our first major assumption: AI as a topic that fits familiar academic categories
- Lesson: the problem you're first asked to solve is real, but it may not be the actual size of the problem

FIRST PRACTICAL RESPONSE

Three Syllabus Positions

Allow

Relatively free use of AI, based on the learning objective

Limit

Permit structured or defined use, with disclosure

Prohibit

No AI use for the assignment or course

The choice depends on the learning objectives, the assignment, the discipline, and the faculty member's judgment: not a single universal rule for every course.

WHAT WE LEARNED

When Enforcement Became the Problem

- A growing number of suspected cases entered the academic integrity system: some valid, many not, many without evidence strong enough to determine what happened
- Faculty were analyzing tone, sentence structure, revision history, and detector scores like an episode of Law and Order
- Inconsistent enforcement: some faculty pursued every suspicion, others stopped enforcing the rule
- False-accusation risk: multilingual students, neurodivergent students, and strong writers could all be misread as suspicious

5×

the normal workload of grading, for a single suspected case

A TURNING POINT

Detection Is Not Evidence

“The AI indicator should not be treated as proof of academic misconduct.”

Turnitin representatives, in conversation with Bhaskar Ghosh and Sean Huss at a University of Florida AI conference

- A detector may be a flag: it is not a factual determination that can carry a disciplinary case
- Any allegation requires human review, corroborating indicators, documented reasoning, and a real opportunity for the student to respond

THE SECOND RESPONSE

Proportionality, Conversation, and Education

- A conversation before a formal charge, when the situation allows it: process, sources, drafts, tools
- First confirmed violation: an instructional consequence, often a zero on the assignment
- Repeated or persistent violations move into the formal academic integrity system
- Faculty must state a policy: they do not have to become adopters, and students are never required to use AI without notice and a reasonable alternative

PROSOCIAL AI USE

Use that supports legitimate learning or work, is transparent when it matters, respects privacy and intellectual property, verifies important claims, and leaves responsibility with the person using the system.

It does not shift risk, deception, or harm onto someone else.

LEARNING WHAT WAS ALREADY HAPPENING

The ATU Survey

634

responses
(485 completed)

65%

already using AI
in their ATU role

~8 in 10

said AI
saved them time

1/3+

somewhat or very
uncomfortable with AI

- ChatGPT was by far the most common tool, followed by Gemini and Copilot
- Use went well beyond writing papers: brainstorming, summarizing, study guides, instructional content, code, rubrics, Grammarly, NotebookLM, Perplexity
- The mismatch that mattered: broad adoption did not equal understanding, confidence, or consistent practice

EXTERNAL BENCHMARKING


The Fragility of a Temporary Task Force

- At a University of Florida AI conference, we met Turnitin, Google, IBM, ethicists, researchers, NVIDIA, and peer faculty from across the country
- Reassuring finding: most institutions were in roughly the same place, with pilots, local guidance, uneven policy, and no mature model sat on a shelf to copy
- Then leadership turnover hit: the AVPAA and the teaching-area chair left just before the conference; Bhaskar left ATU shortly after

Enthusiasm is not institutionalization.

RESTRUCTURING THE WORK

Organized Around the University, Not AI Topics



Angela chairs Information Literacy, Durable Skills & Educational Outcomes. Sean leads Internal Analysis & Effort Effectiveness and central policy drafting. Progress across areas is intentionally uneven.

THE GUIDING PRINCIPLES: PART I

Agree on the Boundaries Before Writing the Rules

- | | |
|--|---|
| <div style="background-color: #2e7d32; color: white; padding: 5px; text-align: center; width: 30px; margin-bottom: 10px;">1</div> Human Stewardship & Accountability | AI assists; it is never an independent institutional decision-maker. Responsibility never transfers to the system. |
| <div style="background-color: #2e7d32; color: white; padding: 5px; text-align: center; width: 30px; margin-bottom: 10px;">2</div> Clear Boundaries | Define what AI may support, what needs approval, what is prohibited. No employee is required to use generative AI. |
| <div style="background-color: #2e7d32; color: white; padding: 5px; text-align: center; width: 30px; margin-bottom: 10px;">3</div> Purpose & Proportionality | The benefit must justify the risk: not adopted because it's fashionable or a vendor added a button. |
| <div style="background-color: #2e7d32; color: white; padding: 5px; text-align: center; width: 30px; margin-bottom: 10px;">4</div> Transparency & Notice | When AI meaningfully influences an outcome, the relevant people should know, and know who owns and is responsible for it. |
| <div style="background-color: #2e7d32; color: white; padding: 5px; text-align: center; width: 30px; margin-bottom: 10px;">5</div> Information Literacy & Verification | If something matters, it must be checked. AI can sound confident while being wrong. |

THE GUIDING PRINCIPLES: PART II

Agentic AI, Ethics, and Continuous Review

6

Privacy & Security

AI tools and workflows are institutional infrastructure the moment they touch university data or systems.

7

Agentic AI

Agents can read, transmit, modify, and act. Any connection to ATU systems needs approval, an owner, limited access, human oversight.

8

Ethical Integrity & Institutional Values

Fairness, nondiscrimination, accessibility, academic integrity, and honest representation still apply.

9

Continuous Evaluation

Approval is not permanent approval. Vendors change features and terms; policy needs review dates and triggers.

AI is never the final authority on: grading • discipline • admissions • financial aid • employment actions • promotion & tenure • access to services

FROM DOCUMENT TO PROCESS

The Policy Factory

Principles

Establish the boundaries

Umbrella Governance Policy

Standing committee, review categories, minimum rules

Domain Modules

Teaching, research, student services, employment, information systems, public-facing operations

Procedures & Practice

Checklists, templates, approved-tools guidance, training

A repeatable process: identify, classify, review, approve or reject, document, monitor, reconsider. Not one document trying to predict every future use.

THE GOVERNANCE SHELL

Committee, Categories, and Review

PROPOSED VOTING MEMBERSHIP

- University counsel · Information security
- Enterprise systems / architecture
- Procurement and contracts
- HR compliance
- Faculty representative · Staff representative

A small standing body by design: the broader Steering Group provides participation; this committee provides formal review and coordination.

THREE REVIEW CATEGORIES

Standard Use

Routine productivity that doesn't influence a high-impact decision or risk sensitive data

Review-Required

Materially influences a high-impact decision, profiles individuals, or uses sensitive data: reviewed before deployment

Prohibited / Exception-Only

Unacceptable residual risk: covert monitoring, fully automated adverse decisions, biometric ID

FIRST MATURE DOMAIN

Teaching and Learning

- Instructor authority preserved: the university sets minimum standards, not a single pedagogy
- Three tiers of use: learning-support, composition & production, evaluation-impacting, each with its own controls
- Grades, feedback, and academic standing remain human decisions; AI may assist, it does not decide
- Detector scores are screening indicators only: cases rest on fabricated citations, contradictions, drafts, or a student's inability to explain their work
- Safe harbor: no one is found in violation because an approved system had a default AI feature they never activated

SECOND MATURE DOMAIN

Cybersecurity, Data, and Agentic AI

- Agentic AI is a new risk category: the system itself can read, transmit, modify, and act inside university infrastructure
- The do-not-paste rule: no sensitive or nonpublic university information into an unapproved AI tool
- Credentials (passwords, MFA codes, API keys, tokens) are never entered into AI tools
- Integrations (email, calendars, files, SIS, HR, finance) are high-risk by default, requiring approval and logging
- Vendor contracts, change control, and code review all apply: approval isn't permanent if the product changes

OPERATIONAL USE

Staff and Administrators

Student Services	Admissions & Aid	HR	Finance & Procurement
MAY Route requests, triage, support communication	MAY Workflows, fraud detection, analysis	MAY Scheduling, drafting, nondispositive workflow	MAY Forecasting, anomaly detection, document review
MAY NOT Deny service, place a hold, impose a sanction	MAY NOT Decide admission, scholarships, or aid	MAY NOT Decide hiring, promotion, discipline, termination	MAY NOT Approve payment or commit the university

FROM COMPLIANCE TO CAPACITY

AI Literacy, Faculty Development, and the Student Course

- Faculty development goes beyond demos: what the systems do, where they fail, what not to enter, how to verify output
- Working with Google on training and tools; Coursera pathways: the credential is not the measure of success, understanding is
- New nontechnical AI literacy course, housed in Behavioral Sciences with English, Library, and CS/Information Systems
- Durable skills, not software training: hallucination detection, source verification, authorship, bias, ethics, human agency

THE EVIDENCE ARM

Internal Analysis and Effort Effectiveness

“Is AI helping ATU, or is it merely changing ATU?”

- Baselines and pre/post testing across training and AI systems
- Measures adoption, trust, learning and work outcomes, equity and access gaps, workload effects, misuse patterns
- Deliverables: campus baselines, AI literacy scorecards, effectiveness and risk reports, an annual State of AI report
- Change is not evidence of improvement. Adoption is not evidence of effectiveness. A vendor demo is not an assessment.

LOOKING AHEAD

Cost, Access, and the Risk of a New Inequality

- Premium AI tools cost real money: the gap between a free tier and a paid tier is not cosmetic
- Same classroom, same assignment: one student works with a stronger, harder-to-detect tool; another is limited to the free tier
- The same pattern repeats at the institutional level: better-funded departments and universities license more capable systems, others don't
- If governance doesn't name this directly, AI risks becoming another axis along which access to opportunity is stratified

WHO CAN AFFORD AI?

- Higher-tier subscriptions put stronger reasoning, and outputs that are harder to flag, behind a paywall
- Free-tier users face real limits, and outputs that are easier for detection systems to flag
- Two people can both “have AI” and not have the same thing

Who owns the system, and who sets the price and terms of access? That has to be an explicit part of our review, not an afterthought.

A DELIBERATE DISSSENT

Are We Adopting AI, or Legitimizing It for Them?

“Any sufficiently advanced technology is indistinguishable from magic.”

Arthur C. Clarke

- When we treat a technology as magic and central to our work, we stop asking who controls it, who pays for it, and what dependencies it creates
- Much of the language of wonder around AI comes from the companies selling access to it, and marketing tends to run ahead of evidence
- Instead of asking what problems genuinely require these tools, institutions often search for problems the tools can solve: adopting and legitimizing vendor products on their behalf
- Once enough time and money is invested, our ability to critically reevaluate the choice shrinks: especially as costs rise and alternatives narrow

TWO WARNINGS FROM OUR OWN CAMPUS

Dependency in Practice

- Institutionally funded access maximizes each person's incentive to use a tool while spreading the cost across the university: an advantage quickly becomes an expectation
- Our Technology Operations Fee rose \$1 this cycle: the largest increase of the round, attributed to software costs. How long will that read as a shared priority if it comes at the expense of raises?
- A free year-long trial of an AI research tool won real faculty enthusiasm, but the ongoing subscription would force cutting a database students actually use, and cutting it later would be far more painful
- Token-based pricing is new and volatile industry-wide in 2026: even the vendors are still finding the floor, and a flat budget cannot absorb an unpredictable, rising bill

This isn't an argument against experimenting: it's a caution that what we fall in love with becomes what we fight to keep funding. Sustainability has to be part of the review, before adoption becomes assumed.

CANDOR ABOUT STATUS

Where We Actually Stand

PRESENTED / OPERATING

- Guiding Principles: presented to the Board of Trustees
- Steering Group and functional areas operating
- Faculty training underway
- AI literacy course developed for a pilot

STILL DRAFT / IN PROGRESS

- Standing governance policy
- Teaching, cybersecurity, and operational modules
- Best-practices guide (needs alignment with newer drafts)
- System inventory and approved-tools process

WHAT WE'D TELL ANOTHER INSTITUTION: PART I

Five Lessons From the First Two Years

- 1 Don't define the problem as student cheating: it's people, purchased systems, and built systems, all at once
- 2 Gather evidence before assuming you know how your campus is using AI
- 3 Establish principles before writing a large collection of rules
- 4 Preserve existing authority structures: build a coordinating shell, not a parallel system
- 5 Distinguish ordinary use from high-impact use, or governance collapses into two extremes

WHAT WE'D TELL ANOTHER INSTITUTION: PART II

Five More Lessons

- 6 Don't confuse detection with proof: a detector score can't carry the weight of a misconduct finding
- 7 Pair policy with literacy: training can't be only a demonstration of features
- 8 Evaluate results: time saved isn't the only outcome that matters
- 9 Design for turnover: a standing structure must own the work, not a few interested people
- 10 Expect revision: the first policy will not be the last policy

The early failures were not wasted effort. They were information.

CONCLUSION

What AI Must Not Replace

What will AI touch?

Probably almost everything we do.

What must it never replace?

Human responsibility for institutional power.

The goal is not to keep AI outside the university: that is no longer realistic. The goal is to ensure that when AI enters the university, the university remains responsible for what happens next.

DISCUSSION

Where Is AI Entering Your Institution Without a Clear Owner?

- Where is AI already inside your systems, before a policy addressed it?
- Who owns the review when a vendor turns on a new feature?
- What would it take to move from a task force to standing governance?

Dr. Sean Huss

Sociology, ATU AI Task Force

Ms. Angela Black

Information Literacy, Durable Skills
& Educational Outcomes

Arkansas Tech University

Thank you.