IT SUMMIT

One Team: Powering Technology Together

THE UNIVERSITY OF ARIZONA

Responsible AI Use

Choosing the right tool, protecting data, and thinking critically

Office of Responsible AI · University of Arizona

This talk: https://tinyurl.com/UA-ITSUMMIT2026

Presented by:
Carlos Lizárraga (clizarraga@arizona.edu) | Rudy Salcido (rsalcido06@arizona.edu)

Tuesday, November 10, 2026

WHY THIS SESSION

Three outcomes for the next 45 minutes

Choose

Pick the right approved UA AI tool for a task.

Protect

Classify your data into UA tiers before you prompt.

Think

Verify, disclose, and keep your own judgment sharp.

Disclaimer: AI assisted presentation (Brainstorming content ideas and slide creation - *Claude Sonnet 5.5* and *Google Gemini*). Reviewed and edited by the author.

FRAMING

AI is here to stay.

It is already in the email, meeting, spreadsheet, and search tools we use. Our job is to assimilate it well, not to ignore it or to trust it blindly.

ROADMAP

Roadmap: six stops

Stop	Topic	Min
1	Approved UA AI tools	9
2	Data tiers: what goes where	8
3	Think, verify, disclose	8
4	The digital AI divide	3
5	Best practices and daily use	5
6	Resources and your first 30 days	5

APPROVED UA TOOLS

Why institutional tools matter

Personal or consumer accounts

Free and paid tools without an enterprise license are approved for **Public data only**. Terms of use may allow training on your inputs.

UA tools with your NetID

Enterprise agreements, UA-managed sign-in, and no training on your inputs. Which data is allowed depends on the tool.

APPROVED UA AI TOOLS

U of A GenAI: the flagship platform

Private by design

Chats are never used to train models. The university will not access your chats or uploaded files.

Open to the UA community

Sign in with NetID at **genai.arizona.edu**. Models run on AWS hardware; no data is sent to OpenAI or Anthropic.

Built for real work

Many models in one interface, side-by-side comparison, file upload, web search, folders with custom prompts, shareable chats.

APPROVED UA AI TOOLS

Ten models, four families

Family	Models on the platform
Anthropic Claude	Haiku 4.5 · Sonnet 5.5 · Opus 5.5
OpenAI GPT 5.6	Luna · Terra · Sol
Open-weight models	Google Gemma 4 31B GPT OSS 120B Meta Llama 4 Maverick 17B
Amazon	Nova Pro

Model lineups rotate as newer versions arrive. Check the platform for the current list.

Which Model Should I Use? — Quick Decision Guide

START: What does my task primarily require?

Coding / Quantitative Analysis	→	Claude Sonnet 5.5 · GPT 5.6 Sol
Long Document (> 50 pages)	→	Llama 4 Maverick · Claude Opus 5.5
Fast Draft / High Volume	→	Claude Haiku 4.5 · GPT 5.6 Luna
Current / Web-Linked Info	→	GPT 5.6 Terra
Multilingual Content	→	Google Gemma 4 31B
Complex Research / Reasoning	→	Claude Opus 5.5 · GPT 5.6 Sol
Image / Figure Analysis	→	Claude Haiku / Sonnet / Opus (multimodal)
Reproducible / Open Source	→	GPT OSS 120B · Llama 4 Maverick

APPROVED UA AI TOOLS

Copilot and Zoom AI Companion: AI inside your daily apps

Microsoft Copilot Enterprise

Inside Word, Excel, PowerPoint, Outlook, and Teams. Email drafts, meeting recaps, Excel formulas, slide outlines. NetID required.

Zoom AI Companion Enterprise

Meeting summaries and follow-ups for Zoom sessions.

Data rule

Both take **Public and Internal** data only. Not for Restricted or regulated data.

APPROVED UA AI TOOLS

Gemini and Gemini Notebook

Google Gemini

Drafting, brainstorming, and summarizing with your **NetID**. UA has an agreement with Google protecting data of users who sign in with their NetID.

Gemini Notebook (formerly NotebookLM)

Source-grounded Q&A with citations over documents you upload: handbooks, policies, research papers.

Data rule

Public and Internal only. The guide lists Gemini; the same limit is applied to Gemini Notebook.

UA's Four Data Tiers

Public

Anyone can see it

Internal

Not public; low risk inside UA

Restricted

Protected by policy or law

Restricted & Regulated

Specific regulatory rules: FERPA, HIPAA, PII, PCI, and others

Less risk

More risk, fewer approved tools

Ask first: which tier is this data?

DATA TIERS

Regulated data needs the most care

FERPA

Student education records. Remove all PII first, delete the chat afterward, do not use AI to store student records.

HIPAA

Not allowed in U of A GenAI or any enterprise tool listed today. Ask University Privacy about the HIPAA-capable option.

IRB and research data

Follow your IRB protocol and data use agreements. Participant data may not be shareable with AI tools.

PII and PCI

No SSNs, card numbers, or government IDs in any AI tool.

Side by Side: Which Tool Is Approved for What

Tool	Sign-in	Public	Internal	Restricted	FERPA*	HIPAA
U of A GenAI	NetID	✓	✓	✓	✓	✗
Microsoft Copilot	UA account	✓	✓	✗	Limited	✗
Gemini / Gemini Notebook	UA account	✓	✓	✗	Limited	✗
Zoom AI Companion	UA account	✓	✓	✗	Limited	✗
Free / paid consumer tools	Personal	✓	✗	✗	✗	✗

* FERPA data only when used appropriately: remove PII first. Source: UA AI Tools and Data Use Guide.

Restricted data: U of A GenAI only. HIPAA data: Soteria only.

Which Tool for Which Task

Public or Internal tasks	Best tool
Draft event copy, newsletters	Copilot / Gemini
Summarize a meeting	Zoom AI Companion
Excel formulas and reports	Copilot
Q&A across a handbook	Gemini Notebook
Brainstorm and outline	Any approved tool
Rewrite for plain language	Copilot / Gemini

Restricted-data tasks	Best tool
Analyze de-identified student feedback	U of A GenAI
Draft templates from advising patterns	U of A GenAI
Work with application files	U of A GenAI, with care
HIPAA-regulated data	Soteria only
SSNs, card numbers, health records	No AI prompt
Admissions or conduct decisions	Not suitable for AI

Confirm current approvals in the AI Tools & Data Use Guide.

DATA TIERS

Substitute and minimize

Before

"Maria G., honors junior, ID 12345678, asked about her appeal..."

After

"Student A, a junior, asked about an appeal of a decision..."

INTERACTION

Vote: green, yellow, or red?

Case A

Rewrite a public registration-deadline notice

Case B

Turn the internal procedures manual into an FAQ

Case C

Paste an advising note with a student name and ID to draft a follow-up

DISCLOSURE

Disclose how you used AI

Why

Trust, academic and research integrity, and accountability. UA guidance says it is good practice to disclose if and how AI was used.

When

Significant drafts, analysis, research, policies, and anything others rely on.

How

One line: tool, task, and what a person verified.

> Sample: "Drafted with [tool] on [date]; reviewed and edited by [name] for accuracy."

CRITICAL THINKING

Critical thinking in the age of AI

Deskilling

Skills you stop practicing fade: writing, analysis, coding, judgment.

Cognitive offloading

Handing routine thinking to a tool. Useful when chosen on purpose.

Cognitive surrender

Accepting the answer without evaluating it. The risk to avoid.

Research signal: in a survey of 319 knowledge workers, higher confidence in GenAI went with less critical thinking (Lee, H. P., et al. *The impact of generative AI on critical thinking: Self-reported reductions in cognitive effort and confidence effects from a survey of knowledge workers.* CHI 2025).

Higher confidence in AI

Less critical thinking

Associated in a survey of knowledge workers

Higher confidence in yourself

More critical thinking

Source: Lee et al., CHI 2025 (n = 319, self-report)

Think first, then ask: write your own rough answer before you prompt.

CRITICAL THINKING

Think first, then ask

- Write your own rough answer or outline first, then compare with the AI.
- Decide which tasks you keep (judgment, ethics, relationships) and which you delegate.
- Ask the model for its assumptions and uncertainty.
- Keep practicing the skills you want to keep.

Hallucination: Fluent, Confident, and Wrong

Under Honors Contract Policy 4.2, students must submit a signed contract to the Dean's office no later than the third Friday of the term. Late contracts incur a $75 fee (Franke Honors Handbook, 2024, p. 18).



Fictional example created for this talk.

Fluent.

Confident.

Wrong.

Polish is not evidence.

CAVEATS

Hallucination: fluent, confident, and wrong

What goes wrong

Invented citations, policies, deadlines, links, and phone numbers. Summaries that drop key conditions.

Verify routine

Check dates, names, numbers, links, and policy statements against the official source. Read like a skeptical editor.

A Quick Verification Routine

- Dates
- Names & titles
- Numbers
- Links & contacts
- Policy wording

1. Compare with the official source
2. Read like a skeptical editor
3. Ask for assumptions and uncertainty

A minute or two of checking is what makes AI use safe.

CAVEATS

Bias: Where It Enters and Where It Shows Up for Us

Data	Model	Output	Decision
Whose data is missing?	What was it optimized for?	Whose norms are the default?	Who is affected?

WHERE IT SHOWS UP FOR US

- Tone and assumptions about students
- Scholarship and outreach wording
- Stereotyped images

Anything that affects people gets human review.

CAVEATS

Bias: where it comes from, where it shows up

Source

Training data reflects historical gaps and stereotypes.

In our work

Tone toward students, outreach and hiring language, image generation, who gets summarized as typical.

Rule

Any output that affects people gets human review before use.

ACCOUNTABILITY

You own what you send.

AI drafts. People decide. Accountability does not transfer to the tool, the vendor, or the prompt.

Responsible AI Use · UA IT Summit 2026

EQUITY

The digital AI divide

Access

Who has an approved tool, a fast connection, and a modern device?

Skills

Who has had time and training to learn effective prompting?

Confidence

Who feels permitted to experiment?

Inclusion

Language, accessibility, and role differences in who benefits.

EQUITY

Narrowing the gap at UA

- NetID gives the whole UA community access to approved tools at no cost to the user.
- Shared training: AI Insights / AI Info Sessions with recordings for those who cannot attend live.
- Share what works: team prompt libraries and shared conversations in U of A GenAI.
- Make time: managers can protect time to learn and experiment.

PRACTICE

Best practices checklist

	Practice
1	Classify your data before you prompt
2	Use approved UA tools with your NetID
3	Substitute and minimize identifiers
4	Think first, then verify against sources
5	Disclose how AI was used
6	Review anything that affects people for bias
7	Keep a personal or team prompt library
8	Stay current: tools and rules change

PRACTICE

Make AI part of daily work

Week 1

Log in to U of A GenAI. Try one low-risk task with public information.

Weeks 2-3

Use the same prompt in two tools. Save what works as a template.

Week 4

Share one prompt with your team and note how you verified it.

RESOURCES

Build your AI fluency

Resource	What it offers
AI Info Sessions	Thursdays 2-3 PM on Zoom; responsibleai.arizona.edu
AI Insights	Tuesday's at noon on Zoom; responsibleai.arizona.edu
UA Google AI Accelerator	Free Coursera-based learning program for UA members; tinyurl.com/UA-Google-AI-Cert
LinkedIn Learning	Self-paced AI courses for staff
UA Library and UCATT	Visit ucatt.arizona.edu, lib.arizona.edu

RESOURCES

Support for research

Arizona Institute for AI and Society

Consultations, Innovation Pilot Projects, Grants collaborations

Office of Responsible AI

AI Consultations, office hours, and the AI Tools and Data Use Guide.

Research Computing

Platforms such as CyVerse and AI Verde / AI Assistant

TAKE-HOME

Your first 30 days

- Log in to U of A GenAI with your NetID.
- Run one real task in two tools. Compare.
- Classify your data first, every time.
- Join one AI Info Session.
- Tell a colleague what you learned.

THANK YOU

Questions?

responsibleai.arizona.edu · genai.arizona.edu · AI Tools and Data Use Guide · AI Consultations and office hours

Carlos Lizárraga (clizarraga@arizona.edu), Arizona Institute for AI & Society