Cy³ Security™

HBCU FACULTY AI RESOURCE PACKET

Plain-Language AI Vocabulary

25 terms you will hear in meetings about AI, written for every discipline. No computer science background assumed. If you can explain these, you can hold your own in any conversation about AI on your campus.

THE BASICS

TERM	WHAT IT ACTUALLY MEANS
Artificial intelligence	An umbrella term for software that performs tasks we used to think required a person. It is not one technology. Most of what people mean today is the next entry.
Generative AI	Software that produces new text, images, audio, or video rather than just sorting or retrieving existing material. ChatGPT, Claude, Copilot, and Gemini are examples.
Large language model	The engine behind a chatbot. It was trained on an enormous amount of text and predicts what words are likely to come next. That is the whole trick. It is prediction, not knowledge, and that explains most of its failures.
Prompt	Whatever you type in. The quality of the answer depends heavily on the quality of the prompt, which is why prompting is now a teachable skill.
Training data	The material the model learned from. If that material underrepresented your students or your field, the model's output will too.
Model	The trained system itself, as distinct from the website or app you use to reach it. One model can sit behind many products.

WHAT GOES WRONG

TERM	WHAT IT ACTUALLY MEANS
Hallucination	When the model states something false with complete confidence. It invents citations, court cases, statistics, and quotations. This is not a bug being fixed soon. It follows directly from the fact that the system predicts plausible text rather than retrieving verified fact.
Bias	Systematic skew in output that reflects patterns in the training data. It shows up in whose names are assumed to belong to whom, which dialects are marked as errors, and which images a request for “a professor” returns.
Drift	The same prompt producing different results over time as the vendor updates the model. An assignment that worked last semester may behave differently this one.
Overreliance	The human stops checking because the output looks polished. This is the most common and most expensive failure in practice, and it is a human problem rather than a technical one.
Black box	The inability to fully explain why a model produced a particular output. It matters when a decision affects a person's admission, aid, grade, or employment.
AI detector	Software that claims to identify AI-written text. Unreliable, and documented to flag writing by multilingual students at higher rates. Treat a result as a conversation starter, never as evidence in an academic integrity case.
Prompt injection	Hidden instructions buried in a document, a web page, or an email that hijack an AI tool when it reads them. Distinct from a hallucination: this is an attack, not an error, and it matters most for agents that can take actions.

USING IT

TERM	WHAT IT ACTUALLY MEANS
Prompt engineering	Writing instructions that reliably get useful output. In practice it means giving context, stating the audience, showing an example, and asking for a specific format.

TERM	WHAT IT ACTUALLY MEANS
Context window	How much text the model can hold in mind at once. Exceed it and the earliest material quietly drops out of the conversation.
Fine-tuning	Additional training on a narrower set of material so the model performs better on a specific task or in a specific voice.
Retrieval augmented generation	Connecting a model to a specific set of documents so its answers are drawn from those sources. This is how institutions build an assistant that answers from their own handbook rather than from the open internet.
Agent	An AI system given the ability to take actions, not just produce text. It can browse, run code, send messages, or update records. Permissions matter enormously here.
Open weights	A model whose underlying parameters are published, so it can be run on your own equipment rather than on a vendor's servers. Relevant when data cannot leave your campus.

OVERSIGHT AND COMPLIANCE

TERM	WHAT IT ACTUALLY MEANS
AI governance	The rules, roles, and review steps that decide which AI tools your institution permits, for what, and who is accountable when something goes wrong. The syllabus policy is one small piece of it.
NIST AI RMF	A voluntary framework from the National Institute of Standards and Technology for identifying and managing AI risk. It is the most commonly cited reference point in higher education AI policy, and it is readable.
Human in the loop	A requirement that a person reviews and approves before an AI-influenced decision takes effect. The standard safeguard for anything affecting a student's record or standing.
Data residency	Where your data physically sits and under whose law. The question to ask before a tool touches student records.
Model training opt-out	A setting or contract term stating that what your people type will not be used to train the vendor's model. Tools available at no cost often lack this option. Paid institutional agreements usually offer it.
Shadow AI	Staff and faculty using AI tools the institution never approved, usually with the best intentions. Every campus has it. Most have not measured it.

Drawn from *AI for Everyone: No One Is Left Behind* by Ahmad Austin.
Print and share this freely with your campus. Please keep this page intact.
Need help applying it on your campus? aaustin@cy3security.com · ai.cy3security.com

v1.0 · October 2026 · Cy3 Security

Resources at no cost