

A FIELD GUIDE · 2026

Prompt to Publish

A field guide to AI for university communicators. What these tools are, how to direct them, how to verify them, and how to know when to put them down.

RAMI DAUD · UNIVERSITY COMMUNICATIONS AND MARKETING

If you read nothing else

The whole guidebook in one spread. Copy it. Pin it up.

Build every prompt from four parts

Role	who the AI writes as
Context	what it cannot guess
Task	the exact deliverable
Constraints	the boundaries and the do-nots

Verify in three steps, every time

1	**Source check**	no openable source, no confirmed claim
2	**Claim isolation**	highlight every number, name, date, title, quote
3	**Independent confirmation**	check against sources the AI did not give you

The more confident it sounds, the more it deserves checking. Polish is not proof.

The eight principles

1. Use AI to support judgment, not replace it.
2. Never confuse fluency with accuracy.
3. Protect private and personally identifying information.
4. Do not fabricate institutional reality.
5. Verify claims according to their potential harm.
6. Disclose synthetic content when audiences could be misled.
7. Preserve authentic voice and representation.
8. The person publishing the work owns the result.

Direction going in. Doubt coming out. Everything in between is technique.

THE WHOLE METHOD IN NINE WORDS

How to use this guidebook

This guidebook works three ways. Read it front to back for the full argument. Jump to Part 3 when you need the craft: the Blueprint, the gallery of worked prompts, the verification protocol. Or pull the appendices, which are built to stand alone as one-page references.

Margin definitions appear throughout wherever an unfamiliar term first earns its place, and everything defined in the margins is collected again in the glossary at the back.

Contents follow. Terms in gold appear in the glossary.

CONTENTS

1 **Understanding AI** 5

2 **Why this module** 10

3 **The craft** 14

4 **The workshop** 23

5 **Judgment** 27

6 **Building a policy** 33

\+ **Conclusion · Glossary · References · Appendices** 36

margin definitions look like this. They sit beside the text that first needs them, and they all reappear in the glossary.

QUICK-START SPREAD · PAGES 2 AND 3

This spread is built to photocopy, with the Blueprint, the protocol, and the eight principles in one opening.



PART ONE

1

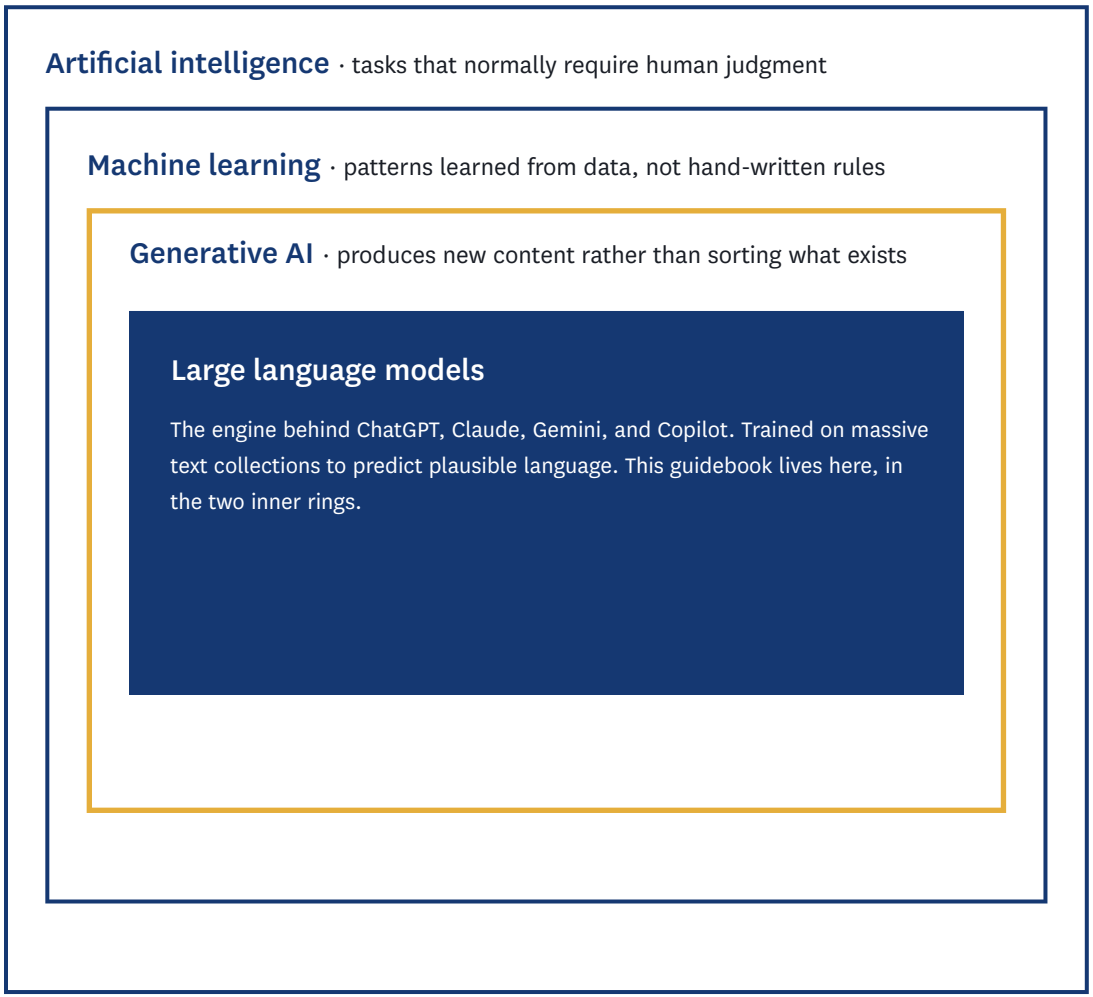
Understanding AI

You cannot calibrate trust in a tool you cannot describe. Start here.

What these tools actually are

Artificial intelligence is not one thing. It is a family of techniques for getting computers to do tasks that normally require human judgment.

The AI in your phone's autocorrect, the AI that flags a fraudulent charge, and the AI drafting your story are different technologies with different failure modes. This guidebook concerns the newest branch: generative AI, systems that produce new text, images, audio, and video rather than sorting or scoring what already exists.



The system that writes your draft has no concept of true. It has a concept of likely.

large language model, LLM *(n.)* a system trained on massive text collections to predict plausible language. The engine behind ChatGPT, Claude, Gemini, and Copilot.

generative AI *(n.)* systems that produce new content (text, images, audio, video) rather than classifying or ranking existing content.

The mental model that matters

PREDICTION, NOT RETRIEVAL · EVERYTHING DOWNSTREAM DEPENDS ON THIS

01

A large language model is a prediction engine. Trained on enormous amounts of text, it learns the statistical patterns of language, then generates the most plausible next words given everything that came before.

02

It does not inherently retrieve verified facts from a reference database. It generates from a compressed statistical impression of its training data, unless it is connected to search, uploaded documents, or another retrieval system.

03

This is why the same system can draft a beautiful paragraph and invent a citation with identical confidence. Fluency and accuracy are produced by the same mechanism, which is exactly why one cannot vouch for the other.

FLUENCY

ACCURACY

how it sounds **vs** **whether it is so**

Two different properties, produced by one mechanism. Nothing about the first guarantees the second.

Seventy years in one page

The overnight revolution was seventy years in the making, with two winters in the middle.

THE LONG RUN-UP

1950–1997

1950

Alan Turing asks whether machines can think and proposes the imitation game

1956

The Dartmouth workshop coins the term artificial intelligence

1966

ELIZA, the first chatbot, fools users with pat-ternmatched therapy

1970s

scriptsThe first AI winter: promises outrun hardware, funding collapses

1980s

Expert systems boom in industry, then bust: the second winter

1997

IBM's Deep Blue defeats world chess champion Garry Kasparov

THE ACCELERATION

2012–2026

2012

AlexNet ignites the deep learning era: neural networks finally have the data and compute they need

2016

AlphaGo defeats Lee Sedol at Go, a game long thought safe from machines

2017

The transformer architecture is published; it becomes the T in GPT

2022

ChatGPT reaches 100 million users in two months, a record for consumer applications at the time

2024–26

Multimodal systems, AI Overviews, and early agentic AI move the question from what can it make to what should we let it do

Two lessons hide in this timeline.

Capability arrives in sudden jumps after long quiet periods, and every boom has carried predictions that aged badly. Hold both.

transformer *(n.)* the 2017 neural network architecture behind modern language models. Its attention mechanism lets models weigh every word against every other word.

What comes next, honestly

Predictions about AI age like milk. Open questions age like wine. Here are the tensions worth tracking, framed as questions because that is what they are.

Agentic AI

Systems that take multi-step actions rather than answering single prompts. The verification question compounds: who checks work that was never seen by a person mid-stream?

For instance: an assistant asked to prep a story package drafts the pitch, assembles the source list, and files the photo request as one chain, with a person reviewing only the end product.

Search is being rewritten

Generative Engine Optimization is emerging alongside SEO as AI Overviews absorb search traffic. Adult learners already trust institutions cited in AI Overviews more, not less (UPCEA, 2025).

For instance: a parent asks a chatbot which Ohio universities offer a program and reads a threesentence answer; whether Kent State appears depends on how clearly our pages state it.

Regulation is unsettled

As of 2026, the EU regulates AI through a risk-based framework while the prevailing US federal posture emphasizes acceleration (European Parliament, 2024; White House, 2025). Institutional policy cannot wait for the law to settle, which is why Part 6 exists.

The authenticity counterweight

Provenance is becoming a technical discipline: edith-istory logging, metadata and rights signals, file-level tagging. These signals declare intent; not every platform will honor them.

For instance: newer professional camera bodies can attach a tamper-evident capture and edit log to every frame, turning authenticity into a documentable property.

Content creation is decentralizing

Faculty and departments can now generate podcasts, videos, and articles without contacting a communications office. Do communicators lead with standards, or inherit the cleanup?

For instance: off-the-shelf tools can turn a research paper into a passable two-voice podcast episode in an afternoon, no studio and no comms office involved.

GEO, generative engine optimization *(n.)* shaping content so AI systems cite and summarize it accurately, the emerging sibling of SEO.

agentic AI *(n.)* systems that plan and execute multi-step tasks with limited human intervention between steps.

Lead with standards, or inherit the cleanup.



PART TWO

2

Why this module

The tools arrived before the training did.

University communications teams adopted generative AI quickly, unevenly, and without a shared playbook. At Kent State's UCM, as at peer institutions everywhere, staff use AI daily to draft stories, caption images, and repurpose content, while the standards governing that use still live mostly in individual judgment.

THREE CONDITIONS MAKE THIS URGENT

- **Adoption runs ahead of governance**

  AI adoption runs ahead of AI governance. Verification practices vary desk to desk.

- **University content carries institutional liability**

  A hallucinated statistic in an enrollment story is a different kind of error than one in a dinner recipe.

- **Trust grows through structure**

  Trust in AI grows through structured, hands-on development, not unstructured exposure (Nazaretsky et al., 2022; Fernández-Batanero et al., 2022).

Standards that live in individual judgment are not standards yet. That is the gap this module closes.

PRACTITIONER CONTEXT

This guidebook was shaped by the author's participation in a professional panel on AI and university photography at the UPAA Symposium, alongside colleagues from Michigan, Missouri S&T, Arizona State, and commercial freelance practice.

Those conversations surfaced the field's live concerns: authenticity, job displacement, disclosure, lagging policy, copyright, and the continued need for human judgment.

It also takes a cue from peer institutions. Arizona State University, represented on the UPAA panel, has embedded this thinking at the institutional level: its computing school is named the School of Computing and Augmented Intelligence, and its president has publicly called for training faculty in the use, concepts, and ethics of AI. The framing positions AI as an extension of human capability rather than a replacement for it. This guidebook takes that framing seriously and gives it operational teeth.

What practitioners are asking

Will AI support our work, or be used to eliminate positions?

How do we protect authenticity in stories about real students and faculty?

When does AI use require disclosure?

Who is accountable when AI-generated content is wrong?

A NOTE ON PROVENANCE

In the spirit of the disclosure standards this guidebook teaches: the Prompt Architecture Blueprint began as a working tool the author developed for UCM colleagues, and this guidebook is its first formalization, a step a division vice president had already inquired about before the project began. Portions of this guidebook were drafted with AI assistance, then verified, restructured, and rewritten by the author. The one AI-generated image inside is labeled as such, deliberately, where synthetic imagery is discussed.

Four things this guidebook teaches

1

Construct

Structured prompts using the four-part Prompt Architecture Blueprint that produce more usable first drafts, measured by a simple before-and-after revision count.

2

Evaluate

AI-generated content for factual accuracy using a three-step verification protocol before anything reaches publication.

3

Identify

The credibility illusion by comparing responses across neutral, directive, and emotionally framed prompts.

4

Assess

Ethical, legal, authenticity, and disclosure risks, and determine when AI should assist the work, when human review is required, and when AI should not be used.



PART THREE

3

The craft

Direction going in. Doubt coming out. Everything in between is technique.

Prediction, not retrieval

THE WORKING MODEL

- These systems predict plausible language; they do not retrieve verified facts.
- Fluency and accuracy are produced independently. Confidence is not accuracy.
- Hallucination is not a simple defect that removes with one patch. It arises from the same predictive mechanism that produces fluent language, though retrieval, stronger models, and verification can reduce the risk.

calibrated trust *(n.)* trust proportioned to evidence and risk. The corrective to both overtrust (it sounds right, so it is) and undertrust (it is dangerous, so never).

The goal is not maximum trust or maximum skepticism. It is calibrated trust: use AI where it creates real value, verify according to risk, and decline to use it when authenticity, privacy, or institutional credibility would be compromised.

THE TWO FAILURE MODES

Overtrust
It sounds right, so it is.

Undertrust
It is dangerous, so never.

The Prompt Architecture Blueprint

Four parts, every prompt. The framework doubles as UCM's standardized prompt intake form.

ROLE

WHO IS THE AI WRITING AS?

A role sets register, vocabulary, and audience assumptions in one move.

CONTEXT

WHAT CAN IT NOT GUESS?

Names, facts, relationships, and materials the model has no way to know. Context is where hallucination gets starved.

TASK

WHAT EXACTLY IS THE DELIVERABLE?

Format, length, components. Vague tasks return vague drafts.

CONSTRAINTS

WHAT ARE THE BOUNDARIES?

Style rules, prohibitions, things it must not invent. Constraints are where institutional risk gets fenced.

The more precise the constraints, the fewer revision cycles the draft needs. Precision up front is also lighter on the planet, a point Part 5 returns to.

The gallery: weak and strong, side by side

Frameworks tell. Examples teach. Five real UCM tasks, each run both ways.

GALLERY 1 OF 5 · THE FEATURE STORY

Read each entry the same way: the weak prompt, what it costs you, the strong rebuild, what changes.

The task: a KSU Today feature on an undergraduate researcher

WEAK

Write a story about the aquatic ecology lab.

What comes back: generic prose, invented atmosphere, a fabricated quote from an unnamed student, no sense of audience. Twenty minutes of cleanup.

STRONG

You are a senior writer for a university news site writing for alumni and prospective families. Context: Kent State's aquatic ecology lab, led by Dr. Costello, hosts undergraduate researchers; the story subject is a junior biology major; the piece runs alongside three photos. Task: a 400-word feature draft with a headline and two subhead options. Constraints: no superlatives about the university, no invented quotes, AP style, do not state research findings the source material does not contain.

What changes: the draft arrives in register, at length, quote-free where quotes were not supplied. The remaining work is voice, not triage.

GALLERY 2 OF 5 · THE CAPTION SET

The task: three Instagram caption options for a photo of the rocketry team's field test

WEAK

Write a fun Instagram caption for this rocket photo.

What comes back: exclamation points, a rocket emoji, and a claim the launch "broke records," which nobody said.

STRONG

Role: social media writer for a university account with a warm, wry voice. Context: photo shows the liquid-fueled rocketry team conducting a cold-weather field test; students pictured have signed releases; no launch occurred, this is instrumentation work. Task: three caption options under 125 characters plus five hashtag suggestions. Constraints: no achievement claims, no emoji, do not name students.

What changes: the model cannot invent a triumph because the constraint forbids claims and the context says what actually happened.

Constraints are where invention gets starved.

GALLERY 3 OF 5 · THE SUMMARY

The task: condensing a 2,000-word grant announcement into a 150-word newsletter item

WEAK

Summarize this press release.

What comes back: a competent summary that quietly promotes the study's hoped-for outcome into a finding.

STRONG

Role: newsletter editor for internal university staff. Context: the attached announcement describes a grant awarded to study X; results do not exist yet. Task: a 150-word item, one quote retained verbatim from the provided text. Constraints: describe the research as beginning, not concluding; every number must appear in the source; flag anything you were unsure about below the draft.

What changes: the constraint that numbers must trace to the source converts the model from author to summarizer. The flag request makes uncertainty visible instead of smoothed over.

Every number must appear in the source.

GALLERY 4 OF 5 · THE SHOT LIST

The task: planning coverage for a college's accreditation-visit story

WEAK

Give me a shot list for an accreditation story.

What comes back: a plausible generic list that misses everything specific to the visit.

STRONG

Role: photo editor planning documentary coverage. Context: three-day site visit; the college requests emphasis on student-faculty interaction in the new simulation lab; two reviewers have declined to be photographed. Task: a shot list of 12 setups ordered by schedule, each with location, subjects, and one-line rationale. Constraints: nothing staged to look like instruction that is not occurring; note where releases are required.

What changes: the no-staging constraint is the authenticity standard from Part 5 written into the prompt itself.

Ethics can be written into the prompt.

GALLERY 5 OF 5 · THE IMAGE GENERATION REQUEST

The task: concept art for an internal presentation, explicitly not for publication

WEAK

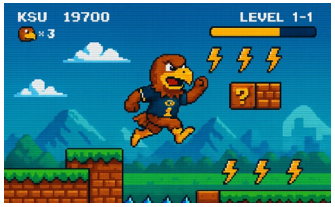
Make an image of diverse students in a science lab.

What comes back: stock-adjacent synthetic people, representation flattened to a template, and an image one forward away from being mistaken for documentation.

STRONG

Role: illustrator creating clearly stylized concept art for an internal slide, not for external use. Context: the slide discusses a proposed lab renovation. Task: one illustration in a flat, graphic, obviously non-photographic style. Constraints: no photorealism, no recognizable individuals, label output as AI-generated in the file name.

What changes: style constraints keep synthetic imagery visibly synthetic. The file-name label makes the disclosure travel with the asset.



The 8-bit illustration on this spread was generated with ChatGPT and is labeled accordingly. That is the standard working: synthetic, stylized, disclosed.

How does AI bluff?

No detector catches everything. Pattern recognition catches a lot.

01

Unverifiable precision

Suspiciously specific statistics with no source attached. Precision is not provenance.

02

The invented citation

Real journal, real-sounding authors, page numbers and all. Always resolve the DOI.

03

Uniform confidence

Human experts hedge unevenly; models hedge evenly or not at all. A draft with no uncertainty anywhere did not weigh any.

04

Studies-show syndrome

"Research consistently shows" followed by nothing checkable.

05

The agreeable mirror

The output adopts your framing, including your errors. Ask for the strongest case against your own premise as a stress test.

06

Stale facts, fresh tone

Confident present-tense claims about roles, prices, and policies that change. Anything time-sensitive gets verified by default.

The credibility illusion

The heart of this guidebook is a finding from original course research (Daud, 2026). Three prompts on the same topic, one neutral, one directive, one emotionally framed, went to the same AI model in separate conversations. The results inverted expectation: the neutral prompt produced the most confident, least-qualified response, the emotionally framed prompt produced the most epistemically careful one, and none of the three cited a single source.

Verification against institutional guidance and survey research surfaced what even the most careful response missed, including 2025 survey data on 760 adult learners showing 56 percent were more likely to trust institutions cited in Google AI Overviews, a finding that complicates the simple erosion-oftrust narrative (Guedon, 2025; UPCEA, 2025).

FIGURE 1 · THREE FRAMINGS, RATED

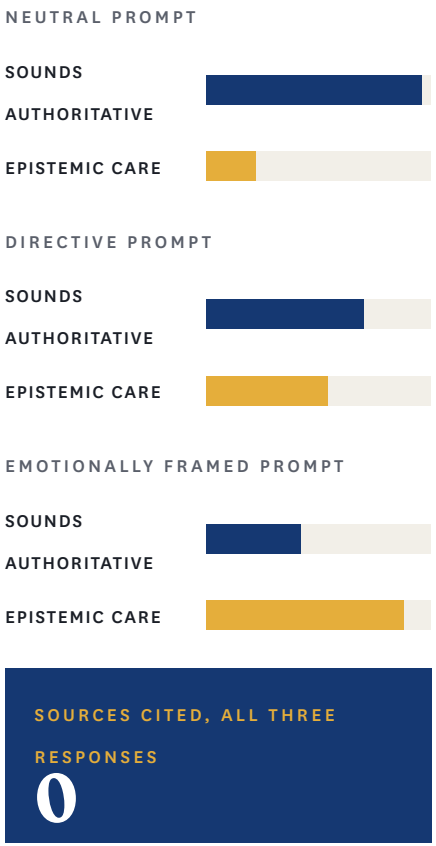


Figure 1. Qualitative ratings from the comparison analysis in Daud (2026). The response that read as most authoritative was the least epistemically careful of the three.

epistemic *(adj.)* relating to knowledge and how carefully it is handled. An epistemically careful response separates what the evidence supports from what remains uncertain.

credibility illusion *(n.)* the gap between how authoritative an AI response sounds and how much epistemic work it actually did. Named in Daud (2026).

The test takes ten minutes with three prompts and one model, and anyone on the team can replicate it.

The finding travels: the method is three conversations and a rubric, not a lab.

The three-step verification protocol

1

Source check

Ask where each factual claim came from. If the AI cannot point to a source you can open, treat the claim as unconfirmed.

2

Claim isolation

Pull every checkable statement out of the draft: numbers, names, dates, titles, rankings, quotes. Highlight them. A claim you have not isolated is a claim you have not checked.

3

Independent confirmation

Confirm against sources that did not come from the AI. Rank them: institutional guidance and survey data outrank commentary. Two AI tools agreeing is not confirmation.

The more authoritative an output sounds, the more verification it deserves. This inverts the the instinct most people bring to polished text.



PART FOUR

4

The workshop

Ninety minutes, hybrid delivery, most of it in participants' hands.

Ninety minutes, block by block

THE HOOK	HOW THESE TOOLS WORK	HANDS-ON 1 · REBUILD YOUR PROMPT	HANDS-ON 2 · SEEDED-ERROR DRILL	SPECTRUM + SCENARIOS	REFLECTION + HANDOFF
0:00	0:10	0:25	0:45	1:00	1:15

Figure 2. Hands-on and discussion blocks, in gold and navy, occupy 65 of 90 minutes.

0:00–0:10 · The hook
Two AI answers, same question. One confident and wrong, one hedged and right. The room votes before knowing which is which. (Objective 3)

0:10–0:25 · How these tools work
Prediction vs. retrieval, calibrated trust, the credibility illusion reveal. (Objective 3)

0:25–0:45 · Hands-on one: rebuild your prompt
Participants bring one current draft or task and rebuild its prompt with the Blueprint, before-and-after comparison. (Objective 1)

0:45–1:00 · Hands-on two: the seeded-error drill
Participants hunt for three planted errors, calibrated obvious, plausible, subtle, in an AI-drafted KSU story. (Objective 2)

1:00–1:15 · The spectrum and scenarios
Small groups place tasks on the risk spectrum, then work one publish-or-don't scenario. (Objective 4)

1:15–1:30 · Reflection and handoff
Exit ticket, one workflow change this week; appendices go home with everyone.

ENGAGEMENT DESIGN

Real materials, not practice topics, a design consistent with facilitation research on adapting professional development to participants' own contexts (Perry & Booth, 2024). Tiered support without labels: a step-by-step handout at every seat, an extension challenge for early finishers.

The AI use spectrum

LOWER-RISK ASSISTANCE

- Brainstorming headlines
- Building shot lists
- Reformatting approved text
- Generating interview questions

MODERATE-RISK ASSISTANCE

- Drafting story copy
- Summarizing source material
- Drafting captions
- Repurposing content

HIGHER-RISK USE

- Generating fictional students
- Fabricating campus scenes
- Altering documentary reality
- Publishing unverified claims

As representational or institutional risk increases, verification, approval, and disclosure requirements must increase with it.

Three scenarios, small groups

The fabricated caption

An AI caption states a student's major, hometown, and program; two details were inferred, not provided. What must be verified? Who owns the correction? Publishable after revision? Disclose?

The synthetic campus image

A department requests a generated photo of diverse students in a lab because no photography exists. Could viewers believe it is real? Does a label fix it? Is synthetic imagery appropriate for recruitment at all?

The efficiency request

A manager proposes AI-generating large volumes of social content to save staff time. What does efficient mean? Which tasks automate safely? Which still require a person?

How learning gets measured

- **Pre/post survey**
- Five items: confidence and knowledge shift.
- **Performance artifact**
- Every participant leaves with one rebuilt prompt and one verified draft from their real workload, plus a before-and-after revision count on that prompt as a simple measure of improvement.
- **Exit ticket**
- One named workflow change within the week.
- **30-day follow-up**
- Actual adoption, feeding the scaling case and a longitudinal research thread.

Every participant leaves with something they built on their own work, not a worksheet.



PART FIVE

5

Judgment

Ethics and evaluation

Privacy and student stories

THREE LAYERS, KEPT HONEST

- **Law**
  - FERPA governs education records and does not always reach marketing content.
- **Policy**
  - Whatever UCM adopts formally, which this guidebook helps build.
- **This guidebook's recommendation**
  - Do not enter student PII, health information, protected academic-record details, or other restricted institutional data into an AI tool unless the use is authorized, the tool is institutionally approved for that data, and any required consent has been obtained. Consent to the story itself remains a separate ethical expectation.

Know which layer you are standing on before anything gets pasted.

Accuracy and institutional liability

- Rankings, research findings, accreditation, outcomes: claims with legal and reputational weight.
- The credibility illusion means the most dangerous errors arrive in the most convincing packaging.

The most dangerous errors arrive in the most convincing packaging.

Disclosure and authenticity are different questions

The proposed working standard: AI-assisted drafting that a human has verified and rewritten needs no public disclosure. AI-generated content published substantially unedited must be disclosed. AI-generated imagery is always disclosed. Offered as a starting position for departmental discussion, not settled policy.

DISCLOSURE ASKS

Have we told the audience how this was made?

AUTHENTICITY ASKS

Even with disclosure, is this an appropriate representation of the institution?

Disclosure is not a magical ethical disinfectant.

- A photograph carries an implied claim that something existed before the camera. University photography often functions as documentation, and prospective students read our images as evidence.
- The field's cautionary examples are real: a university billboard whose diverse student cast attended other schools, and an alumni newsletter that ran a real photograph through AI, cartooned everyone in it, and reached 77,000 inboxes.

PROVENANCE IS BECOMING INFRASTRUCTURE

Cameras that log capture and edit history, file-level metadata and rights signals expressing AI-training preferences, and unit-level tagging standards are turning authenticity from a value into a documentable property, where platforms choose to honor those signals. Communicators who adopt provenance practices early will hold the receipts everyone soon wants.

provenance *(n.)* the documented origin and edit history of a piece of content. It establishes how something was made, not that an image depicts reality truthfully. In photography, increasingly machine-readable.

When not to use AI

The most important AI skill nobody teaches. Some assignments are wrong for these tools, full stop.

Documentary reality

News photography, event coverage, anything presented as a record of what happened. Generation and heavy synthetic editing end the claim to documentation.

Un-consented personal stories

A student's health details, family circumstances, or academic record do not enter a model unless the use is authorized, the tool is approved for that data, and any required consent, including consent to the story itself, has been obtained.

Crisis communications

Speed, pressure, plus institutional stakes, plus hallucination risk is the worst possible combination. AI should not originate or independently determine crisis facts, official positions, or safety instructions; final language must be written or substantively controlled and approved by accountable human officials.

Legal, accreditation, and compliance language

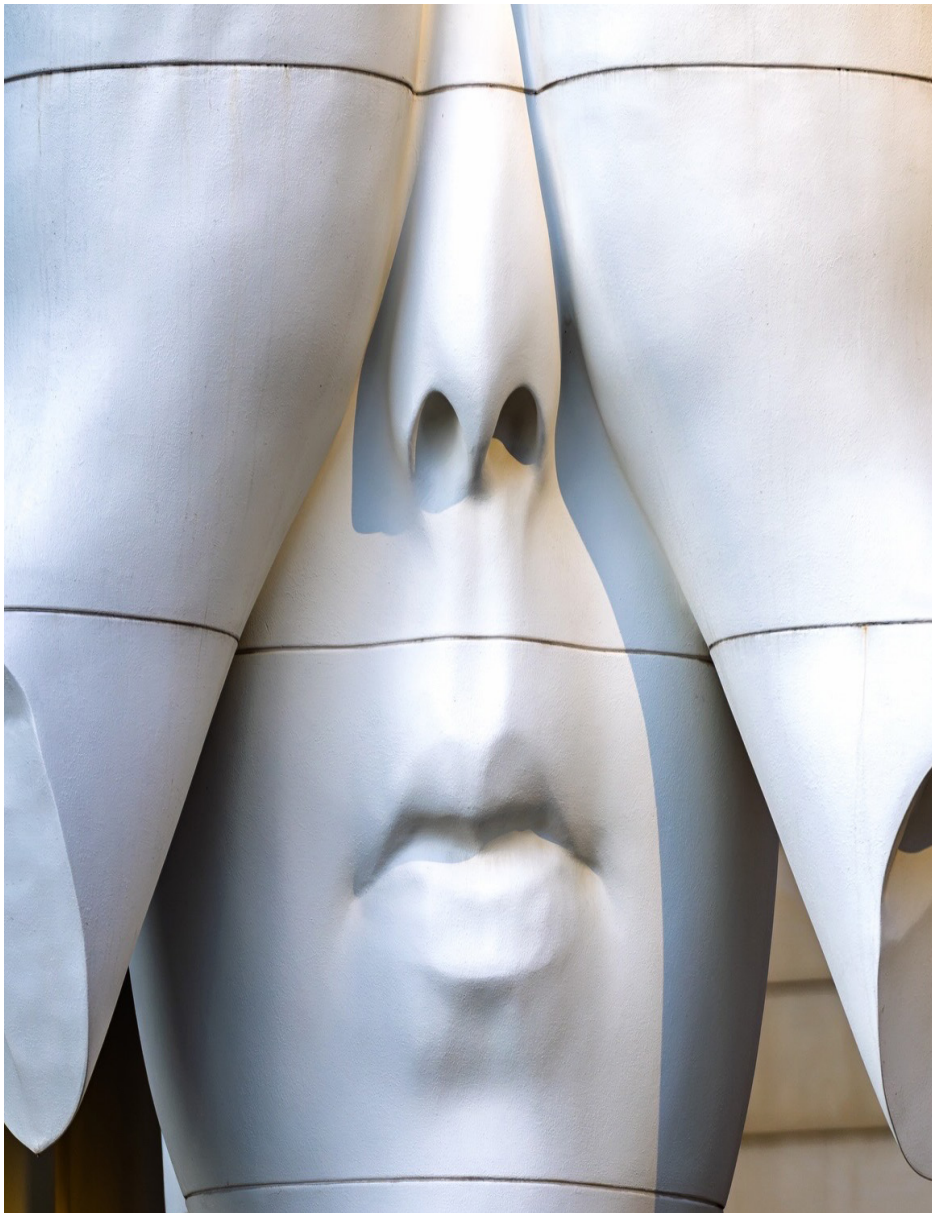
Text with regulatory consequence gets drafted and reviewed by the people accountable for it.

Anything you cannot verify

If no source exists to check a claim against, the claim does not publish. This is not an AI rule. AI just industrialized the temptation.

And the honest cost of saying no

- It would be dishonest to pretend declining is free. A freelancer on the UPAA panel put it plainly: refuse a client's AI request once, and you may never be hired again. At many institutions, pressure to adopt AI is increasingly coming from senior leadership.
- So the skill is not blanket refusal, which is not realistic. It is knowing which uses are worth the professional capital to push back on, and having the verification and authenticity language to make that case credibly.



The judgment call happens at a single desk, where one person decides what publishes.

Refuse a client's AI request once, and you may never be hired again. The skill is knowing which refusals are worth it.

Equity and representation

- Generative systems trained on predominantly Western, English-language data flatten the voices least represented in that data (Birhane, 2020).
- Default AI voice is not neutral voice. Representation is a verification item alongside facts.

Environmental cost

- By common estimates, a generative AI query draws roughly five times the electricity of a conventional web search, and data center cooling consumes about two liters of water per kilowatt-hour (Zewe, 2025). Estimates vary by model and infrastructure; the direction is not in dispute.
- The argument is not against use. It is for intentional use, which the Blueprint already encourages: better outputs in fewer attempts.

5×

The electricity of a web search, per query.

INTENTIONAL PROMPTING IS ENVIRONMENTAL PRACTICE

PART SIX

6

Building a policy

This is not a policy. It is the load-bearing questions any good one has to answer.



Start from stated values

Where institutional leadership has articulated an AI stance, anchor the policy to it; a policy that operationalizes leadership's own language is a policy leadership can adopt. Arizona State offers the model, having gone as far as naming its computing school for augmented intelligence. Where no stated position exists yet, and Kent State has not formally articulated one, the policy process itself is the right occasion to request that framing from leadership.

EIGHT QUESTIONS A WORKING POLICY MUST ANSWER

1 **Scope**

Which tools, which teams, which content types? A policy that covers everything governs nothing.

2 **Risk-to-review mapping**

Adopt the spectrum: as representational and institutional risk rises, so do verification, approval, and disclosure. Write the approval path for each tier, with names attached.

3 **Disclosure tiers**

When is AI assistance silent, when is it credited, when is content labeled synthetic? Part 5's working standard is a starting draft.

4 **Data and tool vetting**

Which tools retain pasted content for training? Who maintains the approved list? What never gets pasted anywhere, regardless of tool?

5 **Provenance and metadata standards**

Decide now how imagery is tagged: capture history, edit logging, AI-training optouts. Cheap to adopt early, expensive to retrofit.

6 **Training and literacy**

Policy without training is shelfware. This guidebook's workshop is a ready-made onboarding unit, repeatable per hub.

7 **Review cadence and ownership**

Name an owner. Set a review date measured in months, not years; the tools will lap an annual policy.

8 **The escalation path**

Where does a staff member take a judgment call the policy does not cover? If the answer is nowhere, the policy will be worked around.

WHAT A POLICY CANNOT DO

No document replaces judgment. The spectrum, the protocol, and the principles travel in people, not PDFs. Policy sets the floor; literacy sets the ceiling.

Rolling this out

FORMAT

90-minute hybrid workshop, repeatable per content hub.

PILOT

One hub team (Arts and Design, the author's facilitator assignment), revise, then scale.

MATERIALS

The deck, the handout, and Appendices B, C, and D as standalone one-pagers; the Blueprint template doubles as the standardized prompt intake form a division vice president proposed formalizing.

COST

Minimal incremental cost. Tools are already licensed or free; facilitation and follow-up draw on existing staff time.

FOLLOW-UP

30-day adoption survey, optional office hours.

SERIES DESIGN

Image generation standards and AI-assisted digital asset management as natural second and third installments.

Prompt to publish

The easy version of AI training teaches people where the buttons are. This guidebook aims at something more durable: a shared standard of practice that lets a communications team move faster without publishing carelessly, built on a framework its members already recognize and a research finding they can test themselves in ten minutes.

The professionals in this field are not waiting for another demonstration of what AI can generate. They are asking what responsible use looks like inside institutions whose credibility depends on truthful representation. The answer this guidebook gives is that prompting is only the beginning. Verification, disclosure, privacy, authenticity, provenance, and human accountability are the competencies that carry an AI-assisted draft safely from prompt to publish.

AI can help us work faster. Our responsibility is to ensure that speed does not outrun truth, authenticity, or judgment.

Glossary

agentic AI · systems that plan and execute multi-step tasks with limited human intervention between steps

AI Overviews · generated summaries at the top of search results, increasingly the first thing prospects read about an institution

artificial intelligence · the family of techniques for performing tasks that normally require human judgment

calibrated trust · trust proportioned to evidence and risk; the corrective to overtrust and undertrust alike

credibility illusion · the gap between how authoritative an AI response sounds and how much epistemic work it did (Daud, 2026)

epistemic · relating to knowledge and how carefully it is handled

fine-tuning · additional training that specializes a general model for a narrower task or voice

generative AI · systems that produce new text, images, audio, or video rather than classifying existing content

GEO · generative engine optimization; shaping content so AI systems cite and summarize it accurately

hallucination · a fluent, plausible statement generated without factual basis; the system has no concept of true

large language model (LLM) · a system trained on massive text collections to predict plausible language

machine learning · techniques that learn patterns from data rather than following hand-written rules

multimodal · a model that works across text, images, audio, and video in one system

prompt framing · the tone and stance a request carries; framing changes the epistemics of what comes back

provenance · the documented origin and edit history of content; in photography, increasingly machine-readable

token · the word-fragment unit a language model actually reads and predicts

training data · the text and images a model learned its patterns from; the source of both its fluency and its blind spots

transformer · the 2017 architecture behind modern language models

EVERY MARGIN DEFINITION IN THIS GUIDEBOOK REAPPEARS HERE · TERMS MARKED IN GOLD THROUGHOUT

References

Birhane, A. (2020). Algorithmic colonization of Africa. *SCRIPTed, 17*(2), 389–409.

Daud, R. (2026). Prompt framing and the credibility illusion: A comparison analysis. MDJ 61001 Module 6 study, Kent State University.

European Parliament. (2024). *EU Artificial Intelligence Act*.

Fernández-Batanero, J. M., et al. (2022). Digital competences for teacher professional development. *European Journal of Teacher Education, 45*(4), 513–531.

Guedon, C. (2025, October 13). AI search and student trust. Search Influence.

Nazaretsky, T., et al. (2022). Teachers' trust in AI-powered educational technology. *British Journal of Educational Technology, 53*(4), 914–931.

NIST. (2023). *AI Risk Management Framework*. (2024). *Generative AI Profile*.

Perry, E., & Booth, J. (2024). The practices of professional development facilitators. *Professional Development in Education, 50*(1), 144–156.

UNESCO. (2021). *Recommendation on the Ethics of Artificial Intelligence*.

UPCEA. (2025, October 24). Adult learner survey on AI search behavior.

White House. (2025). *America's AI Action Plan*.

Zewe, A. (2025). Explained: Generative AI's environmental impact. MIT News.

Professional context: UPAA Symposium panel, June 2026, Ann Arbor, MI.

A door rather than a dead end.

Appendix index

Each appendix is designed to photocopy and stand alone.

A Workshop slide outline · 18 slides

B Verification protocol one-pager

C Prompt Architecture Blueprint template / intake form

D Responsible AI principles

Prompt to Publish

RAMI DAUD · 2026