

ES Designs · Companion resource for the course “AI for Higher Ed Faculty”

AI Tools for Faculty — The Living List

A quarterly-updated guide to the categories of AI tools, the current picks worth knowing, and a portable test for sizing up any tool.

Last full review: May 2026 · Next review: August 2026

Start here (60 seconds)

The course teaches you the parts that don't change: the **categories** of AI tools and a **5-question test** for sizing up any tool. This page is the part that *does* change — the specific tools worth knowing about right now.

Tools get renamed, repriced, bought, and out-classed every few months. So instead of memorizing names, use this page the way you'd use a restaurant guide: skim the category you need, see what's good *this quarter*, and check the date on each entry. We re-review the whole page every three months. If an entry's date is old, trust it less.

A promise and a caution. Every tool below is here because there's real evidence it helps faculty — not because a vendor pitched it. And no tool here replaces your judgment. The research behind this course is blunt about it: the goal is **human + AI working together**, with you reviewing what the machine produces. Speed is the tool's job. The final call is yours.

One rule before you paste anything: if student work or student data is involved, only use a tool your institution has approved with a data agreement. Consumer (free)

versions of many tools can train on what you type. When in doubt, strip names and identifying details first, or don't use it.

The 5 questions (your pocket test)

Run any tool — one on this list or one a colleague swears by — through these five. Two of them are **deal-breakers**.

1. **Privacy: where does my data go?** (*Deal-breaker if student data is involved.*) Does it train on your inputs? Is there an institution-approved version with a real data agreement?
2. **Fit: who's it built for?** General-purpose, or tuned for a specific job like feedback or captioning? Consumer, or education/enterprise?
3. **Accuracy: can I trust the output?** Has it been tested? What does it get wrong? (For anything that cites sources, assume it can invent them — verify.)
4. **Equity: who gets left out?** (*Deal-breaker if you'll require students to use it.*) Does the price or the interface shut out some students or colleagues?
5. **Replaceability: what if it disappears?** If the vendor folds or doubles the price tomorrow, how hard is it to switch?

Good answers are *textured* — "great for X, weak on Y, watch out for Z" — not thumbs-up or thumbs-down.

The 8 categories

Use the jump list, or just scroll.

1. [Drafting & content assistants](#)
2. [Research & literature-discovery assistants](#)
3. [Formative feedback & assessment tools](#)

4. [Tutoring partners](#)
5. [Accessibility & multimodal tools](#)
6. [Image & video generation](#)
7. [Agentic / autonomous AI](#)
8. [Custom micro-models \(Custom GPTs, Gems, Projects\)](#)

Faculty Tools Category Map

Eight categories · One portable framework · The picks cycle on the companion page

THE 5-QUESTION EVALUATION FRAMEWORK — apply to any tool in any category

[gating] 1. Data and privacy posture
2. Audience and use case
3. Verifiable accuracy track record

[gating] 4. Equity and access posture
5. Replaceability if the vendor changes

[gating] = unsatisfactory answer disqualifies the tool

1. Drafting Assistants

General conversational AI for course materials, communications, rubrics.

Largest category; most faculty start here.

2. Research Assistants

Semantic literature search and evidence-synthesis tools.

Verifiability is central — fabrication risk.

3. Formative Feedback Tools

Draft and accelerate feedback on student work, with you in the loop.

Formative, not high-stakes; FERPA gating.

4. Tutoring Partners

Student-facing AI tutoring, often embedded in adaptive platforms.

Faculty role: choose tool, shape prompting.

5. Accessibility & Multimodal Tools

Alt text, captioning, plain language — plus vision, audio, and spatial tools.

Goes beyond text-only accessibility.

6. Image / Video Generation

Diagrams, illustrations, short demos for course materials.

Accuracy + copyright weigh heavier.

7. Agentic / Autonomous AI

Multi-step action taken on its own — risk scoring, 24/7 support copilots.

Institution-level; watch automation bias.

8. Custom Micro-Models

Source-constrained assistants you build — Custom GPTs, Gems, Projects, Copilot.

Role-play & Socratic tutoring; privacy first.

Categories and framework stay current. Specific picks cycle quarterly.

Current best-in-class picks per category: esdesigns.org/ai-tools-for-faculty/

The framework is what survives the next vendor cycle. The picks are what cycle.

Figure 1. The eight categories and the portable 5-question framework. The categories and framework stay current; specific tool picks cycle on this page each quarter.

1 Drafting & content assistants

The everyday workhorses — drafting a rubric, rewriting an email, summarizing a reading, planning a lesson. Most faculty start here. Reviewed: May 2026

- **The big general-purpose four** — Claude, ChatGPT, Gemini, Copilot. All four draft, summarize, and rewrite well. Differences are real but smaller than the marketing suggests. Pick one, learn it deeply. For anything touching student data, use your institution's licensed/enterprise version, not the free consumer one.
- **TeacherMatic** — built *only* for education, with 90+ ready-made generators (lesson plans, rubrics, quizzes). In a multi-university pilot with Jisc, faculty reported saving about **two hours a week**, and remote instructors found it especially useful. A good "I don't want to write the prompt myself" option.
- **Google NotebookLM** — a "walled garden" that answers *only* from documents you upload (your readings, lecture notes, syllabi). Because it stays inside your sources, it largely avoids made-up answers. Turns dense material into study guides, quizzes, mind maps, even podcast-style audio. Through Google Workspace for Education, your uploads are protected and **not used to train Google's public models** — a strong privacy posture.
- **MagicSchool.ai** — fast education-specific scaffolding for rubrics and formative assessments; a solid starting draft you then refine.

Watch for: the free tiers of general tools may train on your input. Move to the institutional version before any student data goes in.

2 Research & literature-discovery assistants

For literature reviews and source-finding. Powerful accelerators — and the category where "it invented a source" is the biggest danger. Reviewed: May 2026

- **Elicit** (~\$120/yr) — automates data extraction for systematic reviews; finds papers by meaning, not just keywords. **Reality check from a Cochrane study:** Elicit was far more *precise* than manual searching (41.8% vs. 7.55%) but caught only **39.5% of the relevant**

studies where human-led searches caught **94.5%**. Translation: excellent for a fast first pass and finding "seed" papers; **not** a replacement for a thorough search.

- **Consensus** (~\$120/yr) — citation-backed answers to direct research questions; good for "what's the consensus on X?"
- **Scite** (~\$144/yr) — shows whether later papers *supported, contrasted, or just mentioned* a source, and flags retractions. In a HKUST evaluation it was uniquely able to read *full text*, not just abstracts — useful for checking a source's credibility.
- **Undermind.ai** (~\$192/yr) — an autonomous "co-researcher" that runs iterative deep searches and adjusts as it goes. Built for exhaustive, high-stakes searches where keywords fail.
- **Perplexity** (~\$240/yr) — fast conversational search for quick scans and fact-checks while prepping lectures.

Watch for: these tools can lean on abstracts and miss nuance buried in the full text, and at least one (Scopus AI, in the HKUST study) produced summaries that copied source wording almost verbatim. Treat output as a strong *draft* synthesis, never the final word — and always confirm a cited paper actually exists and says what the tool claims.

3 Formative feedback & assessment tools

Speed up feedback and routine marking — for formative work, with you in the loop. Not for high-stakes final grades on their own. Reviewed: May 2026

- **Graide** — purpose-built for education (not a generic chatbot). It learns from how *you* mark the first few submissions, then applies that pattern across the cohort. In pilots: about **17% less marking time** (median 16.6 min/script) with a **75% exact match** to human grades at one Western Australian university; at Aston and Kent it generated **4.7–10.5× more written feedback** per student, and **78% of students wanted to keep using it**.
- **Keath.ai** — builds a "digital twin" of *your* grading style from as few as **10 graded essays** plus your rubric, so feedback sounds like you, not a generic machine.

The big rule (from the Jisc marking pilot): use these for **formative** feedback — the rich, fast, "do-this-next" kind students feed into their next draft. Keep them out of high-stakes **summative** grading, where opacity and odd machine errors are too risky. Always **human-in-the-loop**: the tool drafts, you review and sign off.

4 Tutoring partners

Student-facing practice and tutoring — including the AI tutor built into this course's platform.

Reviewed: May 2026

- **Adaptive-platform tutors** — the most relevant option for most faculty is the tutor embedded in the learning platform you already use. The newer ones run a quick check of what a student knows, then adapt explanations as the student works, remembering past sticking points and switching between text, conversation, and visuals.
- **Stand-alone tutoring chatbots** — fine for scaffolded practice; your job is choosing the tool and shaping how it's prompted (e.g., guide, don't just hand over answers).

Watch for: the line between a tutoring partner and an *agentic* system (#7) is blurring — many tutors now act on their own. The difference that matters: a tutoring partner supports *learning*; an agentic tutor also *takes action*. Know which you've got.

5 Accessibility & multimodal tools

Alt text, captions, plain-language rewrites — plus the newer "multimodal" tools that work across images, audio, and physical space. Section 8.2 covers the how; this is the dedicated-tools shortlist. **Reviewed: May 2026**

- **General assistants for quick wins** — Claude/ChatGPT/Gemini handle alt text drafts, reading-level adjustment, and plain-language summaries well. Always read the result; you know your content, the tool doesn't.

- **Vizling** — a research-built mobile app that makes graphic novels, complex maps, and scientific diagrams accessible to blind and low-vision readers by interpreting how the pieces relate spatially — something standard screen readers can't do.
- **Multimodal classroom tools (emerging)** — systems like Stanford's ClassMind analyze classroom video to give instructors teaching feedback. Early but worth watching.

Why "multimodal" matters: the field has moved past text-only. Tools that read across vision, audio, and space open up accessibility that captioning alone can't reach.

6 Image & video generation

Diagrams, illustrations, and short demos for course materials. Higher stakes on accuracy, copyright, and ethics than text tools — so weight question #3 (accuracy/verify) heavily.

Reviewed: May 2026

- **Canva AI** — the most faculty-friendly option: auto-designs slide decks, worksheets, and instructional diagrams from a prompt, so you spend time on the idea, not the formatting.
- **General image generators** — Midjourney, DALL·E, Stable Diffusion and others produce illustrations and diagrams. Quality is high; **accuracy and licensing are not guaranteed**. Check anything factual (a labeled diagram can look right and be wrong), and confirm you have the rights to use what you generate.

Watch for: AI images can introduce subtle errors and copyright gray zones. Verify before you put one in front of students.

7 Agentic / autonomous AI

The newest category — and the one most missing from older tool lists. These don't just answer; they act, across several steps, often without a prompt each time. Reviewed: May 2026 · New this version

- **What it looks like** — student-success systems that watch attendance, submissions, and LMS activity to **flag at-risk students and intervene on their own** (nudges, deadline reminders, auto-scheduled tutoring). Support "copilots" that resolve registration, financial-aid, and IT questions 24/7.
- **Tools to know** — these are platforms institutions buy, not apps you download:
- **Element451 (Bolt)** — an AI agent platform for student engagement and enrollment; as of early 2026 it reports tens of millions of AI-powered student journeys across its customer base.
- **Mainstay** — proactive, evidence-based student messaging (chatbot + SMS) that pairs AI with human coaching; one of the longest-running options in this space.
- **Ocelot** — AI communication and "student engagement" assistant widely used for financial-aid and enrollment questions.
- *Also in the mix:* general-purpose agents (the autonomous "agent" modes now appearing in ChatGPT, Gemini, and Copilot) are starting to do multi-step tasks for individual faculty, too.
- **The evidence so far** — an autonomous-agent deployment at Georgia Southern was linked to **2% enrollment growth and ~\$2.4M in added revenue**; Europe's Open Institute of Technology runs an agent that shifts from study assistant to research tool depending on context.
- **Who chooses it** — usually your institution, not an individual instructor. Your role is to understand it, question it, and keep humans in the loop.

⚠ **The risk to name out loud: automation bias.** It's easy to over-trust a system that acts confidently and on its own. These tools touch sensitive student data and make

consequential moves, so the privacy and accuracy questions (and human review) matter *more* here, not less.

8 Custom micro-models (Custom GPTs, Gems, Projects)

Tools you build yourself — a private AI trained on your own course, kept on a tight leash. Every major platform now offers a version of this; the name changes, the idea doesn't.

Reviewed: May 2026 · New this version

- **What it is** — you upload your own materials (syllabus, rubric, a closed set of readings), write clear instructions, and limit what the model will do. The result is a specialized partner that stays inside *your* content instead of the whole internet — which is what keeps it from wandering off into made-up answers.
- **The same tool under four names** — all no-code, all buildable in about 10 minutes:
- **ChatGPT — Custom GPTs.** The most widely shared; you can publish one for colleagues or students.
- **Google — Gemini Gems.** The natural fit if your work lives in Google Workspace (syncs with Drive/Docs).
- **Claude — Projects.** A persistent workspace that remembers your context across conversations; strong with large document sets.
- **Microsoft — Copilot Studio.** Like the others, but can also *take actions* across Microsoft 365 (update SharePoint, send Outlook mail) — the most "agentic" of the four.
- **What faculty use it for** — authentic **role-play** (WGU Labs built a "Customer Service GPT" for retail/healthcare/call-center practice and a "UDL in Action GPT" where education students adapt lesson plans), **Socratic tutoring** that guides without giving answers, and **rubric-aligned feedback**.
- **Where to learn how** — MIT Sloan, Oxford, and Harvard Business Publishing all have public, faculty-oriented how-to guides.

Watch for: because you're uploading your own materials, privacy (question #1) is front and center — use an institution-approved account, and don't upload anything (especially student

work) you wouldn't want retained. Pick the platform your institution already licenses where you can.

How to use this page each quarter

You don't need to read the whole thing every time. A 10-minute quarterly habit:

1. **Check the dates.** Anything reviewed more than ~6 months ago, trust less until it's refreshed.
2. **Re-run the 5 questions on what you already use.** Tools change underneath you; a pick that passed last year may not now. (This often surfaces a switch you didn't know you needed.)
3. **Skim the two newest categories** (#7 agentic, #8 custom models) — that's where the field is moving fastest.
4. **Before you require a tool of students**, confirm privacy (Q1) and equity/access (Q4). Those two are deal-breakers for a reason.

When a colleague asks *"what tool should I use?"* — send them here, not to a single name.

What changed (changelog)

- **v1 — May 2026.** First version. Built from the course's deep-research review. Added two categories beyond the traditional six — **Agentic / autonomous AI** and **Custom micro-models** — to match where the field has actually moved.
-

A note on the evidence

Every claim with a number behind it (time saved, accuracy rates, pilot results) comes from the research compiled for this course — institutional pilots, peer-reviewed studies, and

library evaluations, not vendor marketing. Where a tool is promising but lightly tested, this page says so. If you want the full citations, they live in the course research files; ask and we'll point you to the source for any specific figure.

Maintained by ES Designs. Questions or a tool we should add? Reach out — this list gets better when faculty flag what's working.

Maintained by ES Designs. This is a living document — picks change; the categories and the 5-question framework do not. Always confirm a tool's privacy and access posture before requiring it of students.