

## AI Agents, Browsers, and Assessment

A new generation of AI-powered agents and agentic browsers can now act autonomously inside Blackboard or Canvas on a student's behalf, reading assignment descriptions, generating responses, and submitting work without any human involvement. Unlike ChatGPT (which requires students to copy and paste between tools), these agents can operate directly within the LMS, indistinguishable from a human user in system logs.

Agentic browsers now available to students:

- Perplexity Comet — free for students since September 2025
- OpenAI ChatGPT Atlas — \$20/month, launched October 2025
- Anthropic Claude in Chrome and Google Gemini in Chrome — in testing/development

Complementary to agentic browsers are **agentic coding platforms** (Claude Code, Clawdbot, Codex, Antigravity). Folks are giving their agents high-level tasks and letting them “crunch away,” overnight or all day, creating and directing other agents.

Previous AI cheating required multiple manual steps (copy prompt, paste into ChatGPT, copy output, paste back, format, submit). AI agents eliminate all of that friction. A student can type "Complete all my assignments due today" and walk away. It is not a new capability in principle, but it removes many of the barriers that made cheating at least somewhat effortful (Mills, Substack, Oct. 2025).

## How AI Agents Can Undermine Academic Integrity

At the AAC&U Institute keynote on agentic AI (Jan. 9, 2026), panelists Anna Mills, Michelle Kassorla, Adam Pryor, and Tim Mousel demonstrated agents **taking quizzes**. The agent navigates to a Blackboard or Canvas quiz, reads each question, selects or generates answers using its training data and course content, and submits. In one viral demonstration, Perplexity's Comet completed an entire Coursera assignment in 16 seconds. They also showed agents **writing discussion board posts** and **submitting essays**. Agents can also **complete entire courses**. Because agents can chain tasks, a student could in principle prompt an agent to complete all pending assignments across multiple courses.

Agentic browsers use normal, authenticated browser sessions. Blackboard or Canvas audit logs cannot distinguish agent activity from human activity. The agent is simply mimicking clicks and keystrokes on the legitimate web interface. This is a dilemma for all course types, but a crisis for those delivered online (Watkins, Chronicle of Higher Education, Dec. 2025).

Privacy is an additional concern. UC Davis research found that AI browser extensions collect and share student academic records with third-party sites, raising FERPA concerns. As panelist Michelle Kassorla emphasized at the AAC&U keynote: it is our responsibility to teach students that using agents in the LMS may expose their own protected data.

## Best Available Solutions

Anna Mills draws on Phillip Dawson's "Swiss cheese" model of academic integrity: no single safeguard is sufficient, but layering multiple imperfect measures can create meaningful friction.

### 1. Improve Materials

- **Assess process as well as product.** Ask for chat transcripts, revision histories, drafts, or reflections on the working process. Assignment sheets should not include excessive detail on "what we've done," which agents use as context.
- Replace standard discussion boards with **social annotation tools** (e.g., Perusall, Hypothesis), where students mark up a shared text. **Voicethread** or **recorded video responses** may also provide additional accountability. These are much harder for agents to engage with. (Recommendation from Adam Pryor, AAC&U keynote)
- Incorporate a **Zoom meeting** into online courses to verify that students are engaged.
- Replace reading quizzes with **"plausible correction" exercises**: give students AI-generated text at a C+/B-level with deliberate errors and ask them to find and fix the mistakes. This teaches critical reading and is more difficult for agents to complete reliably. (Pryor)
- For writing assignments, **try the "Improve the AI" approach**: strip the assignment description to basics, feed it to an LLM yourself, and give students the output. Their assignment is to improve it—through direct revision or by documenting their prompting strategy. (Pryor)
- Consider using **AI detection tools** (like [Pangram](#)) as one part of a layered detection strategy. They can never be relied on as "definitive proof" of academic misconduct, but they can be one more data point to evaluate when trying to identify agentic browser outputs.

### 2. Educate Students

- Focus on **intrinsic motivation**. Help students understand why they are doing each assignment and what skills it develops. The student who completes the work because they see its value is the gold standard.
- Warn students about the **personal dangers of AI misuse**. These risks include learning loss, cognitive offloading that can lead to [diminished critical thinking and cognitive functions](#).
- Warn students about **privacy risks**. Agentic browsers that access Canvas on their behalf may collect and share their academic records with third parties, potentially violating FERPA protections that exist to protect them
- Some pedagogy experts argue that instructors should teach students what agents are and what they can do. Most students don't yet understand the full capabilities or risks. This may be a teachable moment for AI literacy, not just a policing issue. Yet foregrounding agent capability risks increasing student uptake.

### 3. Update Policies and Advocate for Systemic Solutions

- Update AI policies in syllabi to address agents specifically. Most existing policies were written for ChatGPT-style tools and don't cover automated browser-based completion of assignments.
- Press Anthology (Blackboard) for action. Mills and Watkins both urge institutions to submit feature requests, raise the issue in contract conversations, and make clear that LMS companies have a responsibility to protect learning integrity. These are multimillion-dollar contracts with renewal leverage.
- Explore browser-blocking tools as one layer, not the solution. The consensus: blocking is like whack-a-mole and students can work around it, but it's still worth trying as added friction.

---

### Key Sources

- Mills, Anna. "[The Time to Reckon with AI Agents in Digital Learning Spaces.](#)" Substack, Oct. 19, 2025.
- Watkins, Marc. "[Can Educators Counter 'Agentic AI'?](#)" Chronicle of Higher Education, Dec. 4, 2025.
- Watkins, Marc. "[An Open Letter to Perplexity AI.](#)" Rhetorica (Substack), Oct. 17, 2025.
- AAC&U Institute on AI, Pedagogy, and the Curriculum. Keynote: "AI That Takes Action." Jan. 9, 2026. Panelists: Mills, Kassorla, Pryor, MouseL.

Prepared by Andrew Harnish in dialogue with Google Gemini, with input from Kenrick Mock.