

AI use cases gaining traction in higher education operations.

A field guide for enrollment, student success, and finance leaders.

The debate over whether to use AI in higher education is largely over. The real questions now: which uses are worth pursuing, how to know they are working, and what it takes to grow beyond a handful of pilots.

CO-PRODUCED BY



OPENING · THE STATE OF PLAY

The debate is over. The real questions are just beginning.

The conversation about AI in higher education has changed. A year ago, most institutions were still debating whether to explore it. That debate is largely over. The question now is not whether to use AI — it is which uses are worth pursuing, how to know they are working, and what it takes to grow beyond a handful of pilots.

This whitepaper comes out of conversations with practitioners past the exploration stage — people building and running AI agents inside enrollment offices, student success operations, housing departments, and finance teams right now. Their work makes one thing clear: the most valuable applications are not the most visible ones. They are not the chatbot on the homepage. They are agents working inside the operation — eliminating hours of manual processing, surfacing data no one had time to pull, and handling repetitive administrative work that was consuming staff capacity.

The six use cases in these pages span enrollment, student success, and finance. Some are already in production at institutions across the country. All represent the kind of targeted, measurable application that produces real results rather than interesting demos.

But a harder question runs alongside every example: what does it take to go from a successful deployment to a coordinated AI strategy? Getting a use case live is one milestone. Knowing which to prioritize, measuring what they deliver, and sequencing the next investment is a different — and harder — challenge.

IN THEIR WORDS

“The pace of capabilities is very, very fast. So fast that you have to evaluate your AI strategy for where we’re going to be six months from now, not today.”

ZACK PERKINS · CEO & CO-FOUNDER,
COLLEGEVINE

FIG. 01 — THE AI SPECTRUM

Chatbot

Answers questions. Reactive — no action taken.

**Copilot**

Assists a human. It suggests; the person decides.

**Agent**

Does the work end-to-end, escalating exceptions.

*These use cases live at the right end of this spectrum. **Lilac** marks the **agent** throughout this guide.*

SIX USE CASES WORTH YOUR ATTENTION

Use cases real schools are using right now.

Not a comprehensive inventory of everything AI can do. A focused look at applications gaining real traction — each addressing a specific, high-friction problem, each producing a return that can be seen and measured.



- 01** Transcript processing & transfer credit.
- 02** Document verification at upload.
- 03** Fraud detection in submitted documents.
- 04** Scholarship matching & award recommendations.
- 05** The student-worker scheduler.
- 06** Natural-language data querying.

01

ENROLLMENT

AI-assisted transcript processing and transfer credit evaluation.

What took two to three weeks now takes minutes — and the time saved ripples through every step that follows.

Transfer credit evaluation has been one of the most labor-intensive processes in enrollment for decades. Staff manually review transcripts, recalculate GPAs, map incoming courses to institutional equivalencies, and route evaluations to faculty — a sequence that routinely takes two to three weeks. That timeline creates downstream bottlenecks in financial aid, advising, and yield communications, and it leaves transfer students waiting on a decision that can determine whether they enroll.


An AI agent handles the extraction, calculation, and initial mapping that consume the bulk of staff time. When a student uploads a transcript, the agent extracts course data, recalculates the GPA, maps courses to equivalencies, and flags anything requiring faculty review — routing it

automatically. Staff still review and approve. What changes is how much of the work arrives done before a human touches it.

CollegeVine forward-deployed engineer Jessica Bohan, who builds these agents for institutions, describes the shift directly: staff who once spent their days on manual processing can instead spend that time reviewing applications and talking to students. The agent handles extraction and calculation; the human handles the judgment call.

The more consequential change is downstream. When evaluations return in hours rather than weeks, financial aid timelines compress, advising conversations happen earlier, and yield communications move forward without waiting on a review queue.

FIG. 02 — TRANSFER CREDIT EVALUATION: THEN VS. NOW



02

ENROLLMENT

AI-powered document verification.


Confirming that every uploaded document is what it claims to be — at the point of upload, not one file at a time.

Every admission cycle involves a checklist of required documents — transcripts, test scores, identification, background checks, personal statements. Confirming that each uploaded file is actually what it claims to be has historically been a manual, staff-driven task handled one file at a time. At volume, it becomes a bottleneck that backs up the entire review process.

An AI agent verifies authenticity at the point of upload. When a student submits a driver's license, the agent confirms it is valid and extracts identifying information automatically — no manual data entry. When test score reports arrive, it pulls the specific data points the institution needs and maps them to the applicant record. Documents that pass move forward; those that raise questions are flagged for staff review with the specific issue identified.

The same capability extends to more complex documents. When an institution has a defined rubric for a personal statement, an agent can evaluate the submission against it and return a preliminary score — giving staff a consistent, documented starting point before human review begins.

The manual effort of confirming and logging each document is largely eliminated. Staff attention shifts from data entry to exception handling — and because the agent extracts information consistently every time, accuracy improves alongside efficiency.

FIG. 03 – VERIFICATION FLOW**THE SHIFT**

Staff move from data entry to exception handling — the cases where something is genuinely off.

03

ENROLLMENT

AI-powered fraud detection in submitted documents.

Detecting manipulated documents the way a trained analyst would — consistently, and at scale.

Document fraud in higher education is more widespread than most institutions acknowledge — from manipulated transcripts to fabricated background checks to applications built on falsified identities. The challenge is not just that fraud exists; it is that detecting it manually, at scale, is nearly impossible. Staff reviewing hundreds of documents under deadline pressure are not equipped to catch sophisticated manipulation, and the consequences range from misallocated financial aid to significant institutional liability.

An agent built for fraud detection approaches documents the way a trained analyst would. It looks for formatting inconsistencies that may indicate manipulation, cross-references the internal logic within a document — flagging a transcript where the reported GPA doesn't match the individual grades — and identifies patterns across documents that should not share characteristics, such as multiple documents from different institutions bearing the same signature.

When it identifies a potential issue, the agent does not reject the application automatically. It flags the specific anomaly and routes the document for human review, giving the staff member a clear, documented starting point for investigation.

Fraud detection that was previously catch-as-catch-can becomes systematic. The agent does not get tired, miss a flag because the queue is full, or let a document through because a deadline is approaching.

WHAT THE AGENT LOOKS FOR

**Formatting anomalies**

Irregular fonts, misaligned text, altered line spacing.

**Internal inconsistencies**

A reported GPA that doesn't match the individual grades listed.

**Cross-document patterns**

The same signature across documents from different institutions.

IT FLAGS, IT DOESN'T REJECT

Every anomaly routes to a human with the specific issue identified — never an automatic denial.

04

ENROLLMENT & FINANCE

AI-powered scholarship matching and award recommendations.


What takes financial aid staff hours of manual eligibility review, an agent completes in seconds — reasoning included.

Matching students with institutional scholarships is more complex than it looks. Financial aid staff must weigh each student's academic record, financial circumstances, intended major, extracurricular involvement, and other characteristics against dozens — sometimes hundreds — of scholarship criteria. Many awards carry nuanced eligibility rules that require careful interpretation, which makes the process slow and hard to scale. During peak admissions periods, manual review can delay award decisions, slow financial aid packaging, and limit how much personalized guidance staff can offer.

An AI agent completes the initial matching process in seconds. It evaluates a student's profile against institutional scholarship criteria, identifies every award the student is eligible for, explains the reasoning behind each recommendation, and surfaces the strongest opportunities for staff review. Financial aid professionals remain responsible for approving awards and handling exceptions. What changes is how much of the analytical work is done before a staff member enters the process.

Institutions can move faster without adding staff capacity. Students learn what they're eligible for sooner, at one of the highest-stakes moments in the enrollment process. And financial aid teams spend less time interpreting eligibility rules and more time advising students and families on affordability — a process that becomes faster, more consistent, and easier to scale.

FIG. 04 — WHAT THE AGENT WEIGHS



THE SHIFT

Staff spend less time interpreting eligibility rules — more time advising on affordability.

05

FINANCE & OPERATIONS

AI-powered student-worker scheduler.

36 hours of scheduling a year, per supervisor, reduced to under two.

Scheduling student workers looks straightforward until you are the one doing it. Supervisors must work around each student's class schedule – which they often cannot access directly due to FERPA restrictions – while complying with labor requirements around maximum hours, required meal breaks, and shift lengths. At institutions with significant student employment programs, supervisors can spend more than 30 hours per year on scheduling alone. It is laborious, low-value work, and it starts over every time a student changes their class schedule.

An institution on the West Coast built a scheduling agent embedded directly in their employee-management system. The agent knows which students each supervisor manages, which courses those students are enrolled in, and the full set of scheduling rules that apply – hours per day, hours per week, break requirements, and standard operating hours. A supervisor opens the tool, inputs any special parameters for the week, and the agent produces a complete proposed schedule that accounts for class conflicts, coverage gaps, and labor compliance.

The more significant shift is qualitative: supervisors stop spending their time on a logistics problem and start spending it on the work that requires their judgment and their relationship with the students they manage.

36 hrs → **2** hrs

PER SUPERVISOR, PER YEAR – SPENT ON SCHEDULING.

1,360 hrs

Returned annually across an institution with 40 student-employment supervisors – redirected toward mentoring, performance, and running the department.

06

ENROLLMENT & FINANCE LEADERSHIP

Natural-language data querying and automated leadership reporting.


When asking costs nothing, leaders ask the questions they should be asking.

Getting data to the people who need it in higher education is slow. When a senior leader needs enrollment metrics for a board meeting, they call the chief of staff, who calls the VP, who calls the director, who pulls the report and sends it back up the chain. By the time it arrives, two or three days and twenty-plus hours of collective staff time have been consumed — often to answer one question that prompts three more. The consequence: leaders stop asking questions they should be asking, because they know what it costs the people below them.

AI-powered analytics is beginning to change this dynamic. Rather than building queries or requesting reports, leaders ask in plain language: show me deposit pace compared to

last year; identify accepted students who have not yet deposited; flag unusual patterns in this cycle's funnel. The agent queries the connected systems, runs the analysis, and returns the result — often as a dashboard it has built on the fly.

The same capability extends to scheduled, automated reporting. An enrollment leader can ask for a deposit summary every Monday morning; a finance leader can request a weekly anomaly report. It arrives without anyone having to produce it. Leaders ask more questions because asking costs nothing — and the team operates with better information, faster.



“The strategy you’re bringing to this is actually the value you create. Freeing you from the work of just getting to the data — that’s where the real change is.” – ZACK PERKINS, COLLEGEVINE


FROM USE CASES TO STRATEGY

A good use case is a milestone. A strategy is the advantage.

The six examples in this whitepaper are not a complete picture of where AI is going in higher education. They are a representative slice of what is working right now — applications built, deployed, and put to work inside real institutional operations. What they share is specificity: each addresses a defined problem, operates within a clear scope, and produces a result that can be measured. That specificity is not incidental. It is the reason they work. The challenge for most institutions is not a shortage of ideas — the opportunity landscape is expanding faster than most teams can track.

A scattered approach to AI, chasing interesting use cases without a framework for prioritizing, resourcing, and measuring them, produces scattered results: some deployments gain traction, others sit underused, and the potential compounds without ever fully delivering. The institutions seeing the most durable returns are not necessarily moving fastest — they are moving deliberately, with a clear picture of where capacity is being consumed and how to sequence investments so early wins build confidence for what comes next. The harder discipline is knowing which place is right for your institution — and building the measurement habits that turn a deployment into a lasting advantage.

FIG. 06 — A LENS FOR SEQUENCING USE CASES



ABOUT THIS FIELD GUIDE

Co-produced by Human Capital Education and CollegeVine.

A practitioner's look at the AI use cases that have moved past the pilot stage — and into the daily operation of real institutions.

CO-WRITTEN BY



SCAN TO LEARN MORE



Human Capital Education



CollegeVine