

Essay *Helper*

The AI writing coach
that *never invents a source*.

A trustworthy writing tutor that lives inside Microsoft Word — it reviews a student's draft, coaches their citations, and leaves feedback in the margin, without ever fabricating a fact or doing the thinking for them.

● Shipped · v1.4

155 tests green

Live model-safety evals passing

Working software, not a concept

THE ONE LINE

*Feedback at the scale of
AI, with the integrity of a
good teacher.*

Every student needs a patient reader who will tell them where their argument is thin and where a claim needs a source. Almost none of them get one in time. Generic AI fills the gap — but it hallucinates facts, invents citations, and quietly writes the essay for them, which is worse than no help at all.

Essay Helper is the alternative built for education specifically: an AI coach that is architecturally incapable of inventing a source, that coaches instead of ghost-writing, and that meets students where they already write. It is built, tested, and shipped today.

THE PROBLEM

*Three failures sit between
a student and honest,
timely feedback.*

Good feedback doesn't scale. Bad AI does.

The tools students actually reach for are the ones least suited to learning.

01

Teachers can't be everywhere

One instructor, thirty drafts, one week. Substantive, per-paragraph feedback on argument and evidence simply doesn't fit the hours available.

02

Generic AI makes things up

Ask a chatbot to check a citation and it will happily invent an author, a year, a page number. In academic work, a fabricated source isn't a bug — it's misconduct.

03

Or it just does the work

The path of least resistance is "write my essay." That trades a grade for the learning, and leaves institutions fighting an integrity crisis instead of teaching.

WHAT IT IS

*It opens in the pane
beside the document and
works on the real draft.*

A coach in the margin of the page

Essay Helper is a Microsoft 365 Word add-in. The student describes the assignment in a short chat, picks how they want to be coached, and hits [Review](#). It reads the whole draft and leaves comments anchored to the exact sentences that need them — the same place a teacher's pen would land.

1

Understand the assignment

A quick conversation captures the prompt, level, and rubric — and whether it's an essay, worksheet, lab report, or reflection.

2

Review the draft

It flags weak arguments, unsupported claims, and citation problems — anchored to the sentence, in the document.

3

Coach, don't ghost-write

Depending on the mode, it asks the guiding question, gives the mechanical fix, or — in Corrections mode — answers directly and labels exactly where the answer came from.

WHERE WE ARE TODAY

0

Sources ever invented
— guaranteed by
construction, not by a
prompt

155 ✓

Automated tests green,
plus live model-safety
evals passing

6

Coaching modes, from
Socratic to direct
answers

5

Assignment types the
review adapts itself to

WHY IT WINS

The guarantee is a property of the code, so it can't be jailbroken away.

The moat is trust — and it's built in

THE CORE GUARANTEE

It never invents a source.

Most AI safety is a polite request in a prompt that a clever user can talk around. Ours isn't. A separate validation layer inspects every citation the model produces and **drops any that don't point to a source the student actually supplied** — before it ever reaches the page. The promise holds even if the model misbehaves.

How it holds: answers live in plain comments; citations travel a separate, validated channel. The model can offer an answer from its own knowledge — clearly tagged "*verify this*" — but it can never manufacture a **source**. That boundary is covered by automated tests and by live evals run against the real model.

Aa Adapts to the assignment

An essay is judged on thesis and evidence; a worksheet on grounded answers; a lab report by its sections. One tool, the right expectations.

Coaches, in six stances

Socratic, mechanics-only, editor, critical, encouraging — and Corrections mode, which answers directly and tags every answer's provenance.

“” Honest by design

It labels an own-knowledge answer versus a grounded one, and every screen carries the reminder: *double-check answers — this tool can be wrong.*

↪ In the flow of writing

No new app to learn. It works in Word, where the draft already lives — comments land on the sentence, feedback survives across sessions.

§ Private by architecture

The AI key and every model call stay on the server. The student's add-in never carries a secret, and never talks to the model directly.

↔ Portable core

The engine is separated from Word by design — the same reviewer can front a Google Docs or web client with no change to the guarantee.

THE FIELD

*Everyone does one job.
Essay Helper is the only
one doing this one.*

No one else coaches *and* stays honest

	ESSAY HELPER	GRAMMARLY	GE
Feedback on argument & evidence	●	○	
Coaches without writing it for them	●	◐	
Never invents sources or facts	●	—	
Works inside the document	●	●	
Respects academic integrity	●	◐	
Available 24/7, at scale	●	●	

● full ◐ partial ○ no — not applicable

WHY NOW

*The demand and the fear
arrived in the same year.*

A market that just became urgent

Generative AI went from novelty to default in classrooms in under two years. Institutions now face the same question at every level: **how do we let students use AI without letting it fabricate, plagiarize, or replace the learning?** The tools on offer are either general-purpose chatbots with none of those guardrails, or plagiarism *detectors* that police students rather than help them.

Essay Helper is positioned exactly in that gap — an AI that administrators can approve *because* of how it's built. The buyer (a department, school, or university) is actively shopping for a safe, aligned writing tool, and there is no incumbent that both coaches and guarantees integrity.

Reach

Hundreds of millions

of students write assignments in Word or Docs every term — the surface the product already runs on.

Timing

Mainstream, now

AI use in coursework is no longer fringe; policy and procurement are catching up fast.

Gap

No safe incumbent

Coaching tools ignore integrity; integrity tools ignore coaching. We do both.

HOW IT EARNS

*Free for the student;
licensed by the institution
that needs the guarantee.*

Business model

Land with individual students; expand to the institutions that need a tool they can sanction. Figures below are an illustrative starting model.

STUDENT • FREE

\$0

-  Core review in Word
-  A few reviews a month
-  Socratic & mechanics coaching

PRO

\$8 / student / mo

-  Unlimited reviews
-  All six coaching modes
-  Corrections & conciseness controls

INSTITUTION

Site license

-  Per-seat or per-FTE pricing
-  Admin controls & SSO
-  Teacher-facing analytics

The institutional license is the revenue engine: it's sold on the one thing individuals can't buy elsewhere — a writing AI the institution can *vouch for*.

WHAT'S NEXT

The hard part — a shipped, safe engine — is already done.

Roadmap



SHIPPED

The product, live DONE

In-Word review, citation coaching, six modes, five assignment types, Corrections mode, conciseness controls, the no-fabrication guarantee, and the honesty disclaimer — all built, tested, and merged.



NEXT

History-aware review

Compare a draft against its earlier versions — praise real revision, catch regressions, and coach the arc of the writing, not just the snapshot.



THEN

Google Docs & teacher analytics

Bring the same guaranteed engine to Docs (the core is already client-agnostic), and give instructors a view of where their class struggles.



LATER

Institutional platform

Rostering & LMS integration, admin dashboards, more citation styles and languages — the surface a site license is sold on.

Risks & mitigations

The model drifts or misbehaves over time.

Contained: the citation guarantee is a code-level validator, not a prompt, and behavior is gated by live evals against the real model before anything ships.

Dependence on the Microsoft Word platform.

Contained: the reviewing engine is deliberately separated from the Word client — the same core can serve Docs or web with the guarantee intact.

Institutions are slow, cautious buyers.

Turned into an edge: caution is exactly why they'll choose the tool built to be safe. We land free with students while the institutional sale matures.

"Isn't this just ChatGPT?"

Answered by the demo: ask both to check a citation. One invents an author; ours refuses and asks for the real one. The difference is the product.

THE ASK

Let's put it in front of real students

The engineering risk is behind us — the product is built, safe, and shipped. What it needs now is a first classroom and a short runway to the next surface. Three things unlock the next stage:

GREENLIGHT

A pilot with 2–3 design-partner classes this term

FOCUS

A short push to Google Docs + teacher analytics

BACKING

Support to court the first institutional buyer

About this document. Every product, architecture, and engineering claim reflects the current, shipped codebase (Essay Helper v1.4): the no-fabrication validation layer, the six coaching modes and five assignment types, the 155 passing automated tests, and the live model-safety evals. **Market sizing, pricing, and financial figures are illustrative planning assumptions** for discussion — not audited data or forecasts. Prepared as an internal business case · Confidential.