

A THROUGHLINE.AI PRODUCT

Starting Plan v0.2

Problem, customers, assumptions, and product strategy for TheDirectLine.AI customer discovery

Prepared by Everrett V. Harris, Founder

October 25, 2026 | Living document for Harvard Innovation Labs and MIT Media Lab mentors

Draft for discussion. Figures cite public sources listed and verified in Section 14. Hypotheses are labeled as hypotheses.

CONTENTS

What is inside

01	Executive summary
02	The problem
03	Value proposition
04	Customer segmentation
05	Pain points
06	Assumptions and how we grade them
07	Customer discovery plan
08	Product: what we build
09	Three approaches to the goal
10	Evidence exhibits
11	Business model and metrics
12	90-day roadmap
13	Risks, mentor asks, and open decisions
14	Sources and verification log
A	Appendix: ThroughLine.AI deliverable standard

HOW TO USE THIS DOCUMENT

Every bracketed number is a live link to its source. Click [1] to open the Pew report. Section 14 lists all 27 sources with the exact figure checked and its verification status as of Oct 25, 2026. Update the assumption scores in Section 06 after every interview batch. Log changes in the version line on the cover.

WHAT CHANGED SINCE v0.1

v0.2 adds 8 charts, 14 new sources, clickable citations, a value proposition section, and a verification log. It removes four figures from v0.1 that could not be traced to a primary source: the 68% urban teacher training figure, the 27-state legislation count, the 80% of students not taught figure, and a secondary income analysis now replaced by the Pew report itself.

SECTION 01

Executive summary

TheDirectLine.AI trains students ages 15 to 22 from marginalized communities to use AI the way power users do. They learn to write strong prompts, judge outputs, iterate, and ship finished artifacts.

CORE HYPOTHESIS

Students from low-income, Black, and Hispanic communities already use AI as much as or more than their peers. Few are taught to use it well. The gap is depth, not access. If we give them structured practice with rubric feedback and real artifacts, their output quality rises measurably, and schools and programs will pay for that result.

Black and Hispanic teens use AI more, but are half as likely to be AI-ready

U.S. teens. Left: share who used chatbots for schoolwork help. Right: share in the highest AI-readiness group.

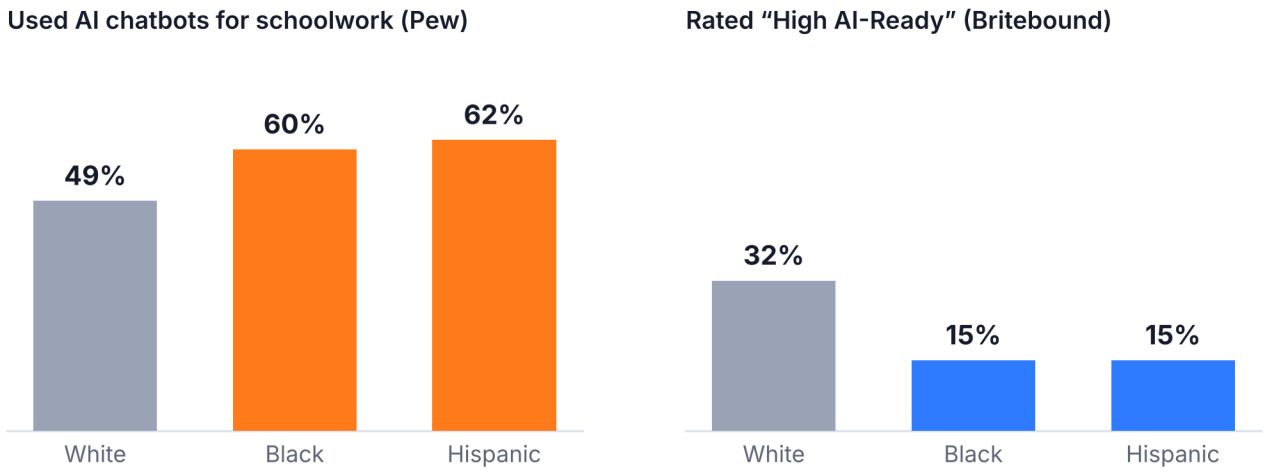


Chart 1. The core value proposition in one view. Use is high in our segment. Readiness is not. Different surveys, shown side by side. Source: [1] Pew Research Center, Feb 24, 2026 [5] Britebound, Jun 2026

The plan in five points

- **Problem.** 60% of Black teens and 62% of Hispanic teens use chatbots for schoolwork [1]. Only 15% of each group rank as highly AI-ready, versus 32% of White teens [5]. Only 36% of teens say school teaches them enough about AI [5].
- **Customer.** Users are learners ages 15 to 22. Our first buyers are program directors at college-access nonprofits and community colleges. School districts are the long-term buyer.
- **Product.** A practice and feedback platform. The loop is brief, prompt, output, check, revise, ship. Every challenge has a rubric, an exemplar, and a verification step.

- **Quality bar.** Zero-shot AI grading agrees poorly with human graders [24] [25]. No challenge ships until our feedback agrees with educator grading on 85% of a 100-attempt eval set, with zero fabricated facts.
- **Next 90 days.** 30 discovery interviews, a 5-challenge MVP built with Claude Code, and one pilot of 30 to 60 learners with blind-scored results.

Recommended approach

Section 09 compares three paths. We recommend leading with a **Program-in-a-Box** license for nonprofits and community colleges. Pilot outcomes from that path become the evidence package for district sales in the 2027 to 2028 school year.

SECTION 02

The problem

AI is now part of how young people do school and how employers screen entry-level talent. Students in our segment use it often, but most were never taught to use it well. Usage without mastery widens the gap AI could close.

PROBLEM STATEMENT

Students ages 15 to 22 from marginalized communities use AI chatbots at equal or higher rates than their peers, but they lack instruction in structuring requests, checking outputs, iterating, and producing finished work. Their schools are the least likely to train teachers. Employers now expect these skills on day one.

Use is rising fast

Teen use of ChatGPT for schoolwork doubled from 13% to 26% between 2023 and 2024 [3]. By fall 2025, 54% of teens used some chatbot for schoolwork, and about 3 in 10 used chatbots daily [1].

Teen use of AI for schoolwork has climbed every year

Share of U.S. teens ages 13 to 17.

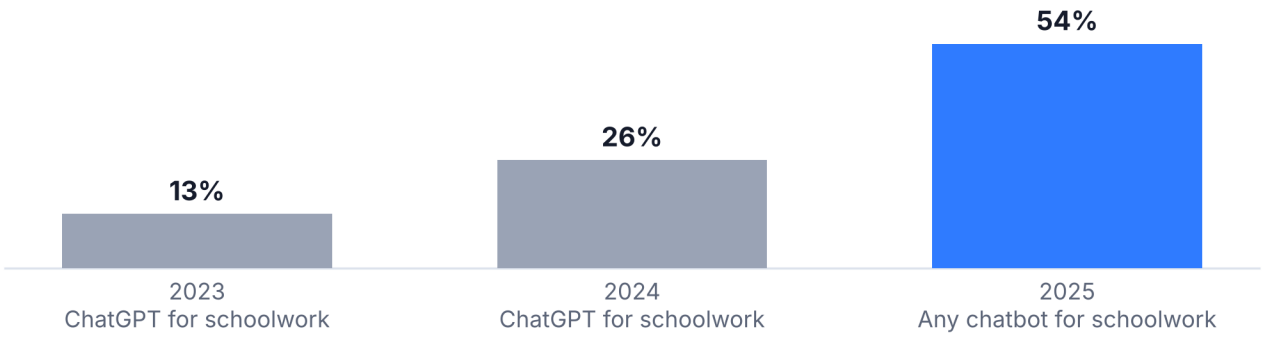


Chart 2. Pew changed the question from ChatGPT to any chatbot in 2025, so the third bar is not a like-for-like trend point. Source: [3] Pew Research Center, Jan 15, 2025 [1] Pew Research Center, Feb 24, 2026

Layer 1: The depth gap

Use is not the problem. Depth is. Black and Hispanic teens are more likely to use chatbots for schoolwork [1] [4], and 18% of Black teens do all or most of their schoolwork with them, versus 6% of White teens [1]. Yet 29% of Black teens limit school AI use to seeking homework answers [5]. First-generation, low-income college students show the same pattern: adoption is real but shallow, limited to chatbots for narrow tasks [6]. Black and Hispanic undergraduates report lower conceptual AI literacy [7].

Answer-seeking only: the shallow use DirectLine is built to fix

Share of teens whose school AI use is limited to seeking homework answers.

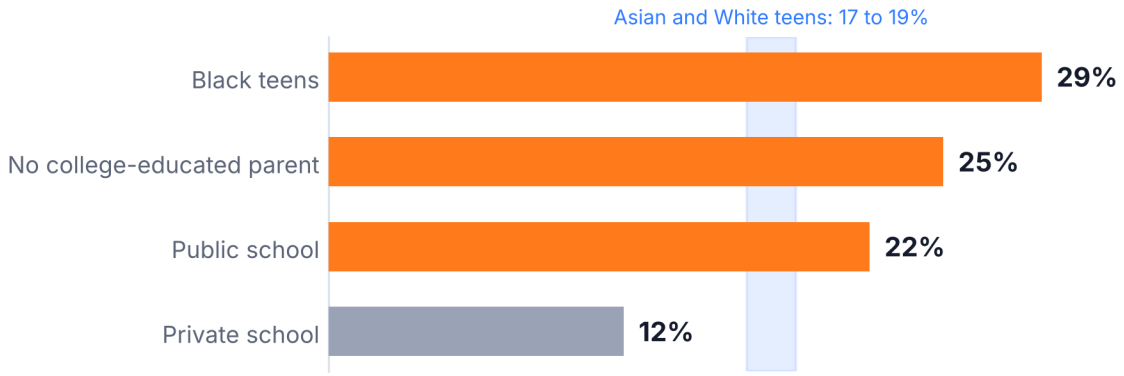


Chart 3. Answer-seeking only means students use AI for homework answers and nothing else. DirectLine moves them to drafting, checking, and building. Source: [5] Britebound, Jun 2026

Lower-income teens lean on AI most, with the least guidance at home

U.S. teens and parents by household income.

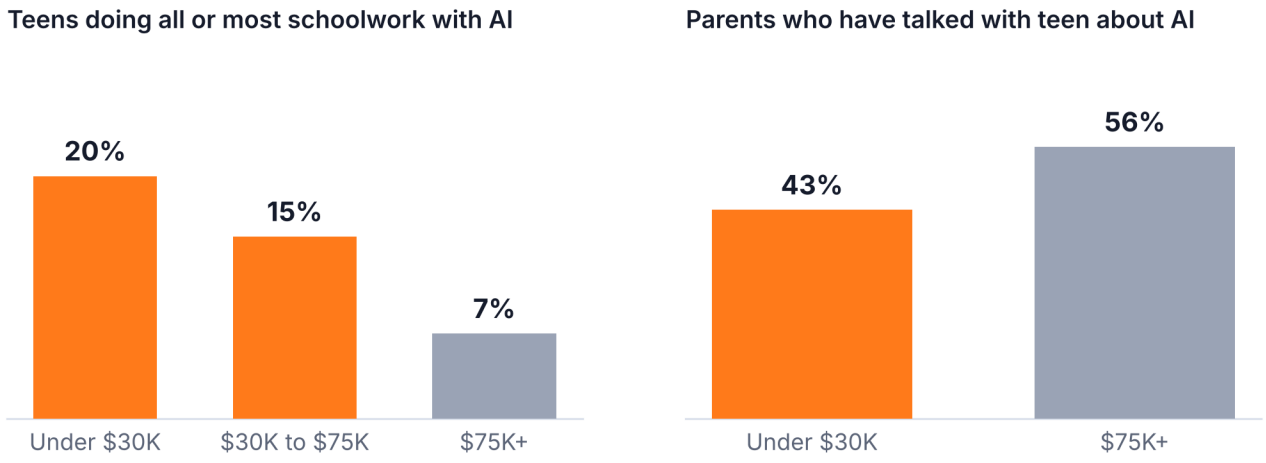


Chart 4. Teens in households under $30K are almost 3x as likely as teens in $75K+ households to do most schoolwork with AI, and their parents are less likely to have discussed AI with them. Source: [1] Pew Research Center, Feb 24, 2026

Layer 2: The instruction gap

78% of teachers say responsible AI use belongs in the curriculum, per Ipsos [9]. The supply of instruction runs the other way. In fall 2024, 67% of low-poverty districts trained teachers on AI, versus 39% of high-poverty districts [8]. Only 13% of public school teens say most teachers encourage AI use, versus 57% at private schools [5]. Districts also struggle to fit AI literacy into the schedule [17].

The students who need instruction most are least likely to get it

Teacher training and classroom support for AI, by school context.

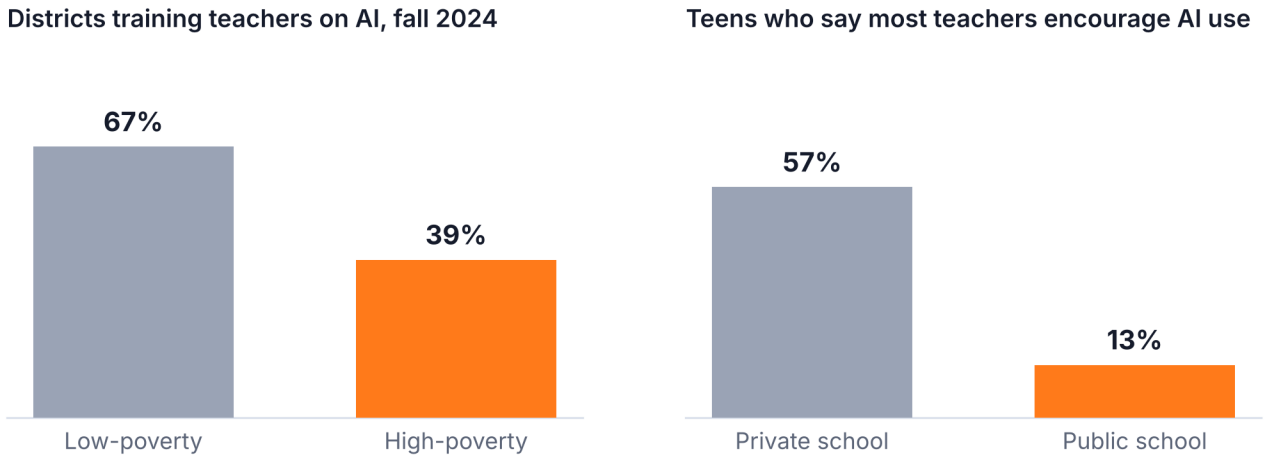


Chart 5. Students in high-poverty districts and public schools get the least AI instruction. Source: [8] K-12 Dive, reporting a RAND survey, Apr 18, 2025 [5] Britebound, Jun 2026

Layer 3: The labor market gap

Employers now pay for AI skill. AI-exposed entry-level roles are 7x more likely to require senior skills like judgment, and those roles grew 35% from 2019 to 2025 while other entry-level roles shrank 10% [10]. Employers say more than one-third of entry-level jobs now require AI skills, nearly triple the share in fall 2025 [11]. Yet 50.5% of college students say they are not building AI skills [12].

Employers now pay for AI skill and judgment at entry level

Demand signals from job postings and employer surveys.

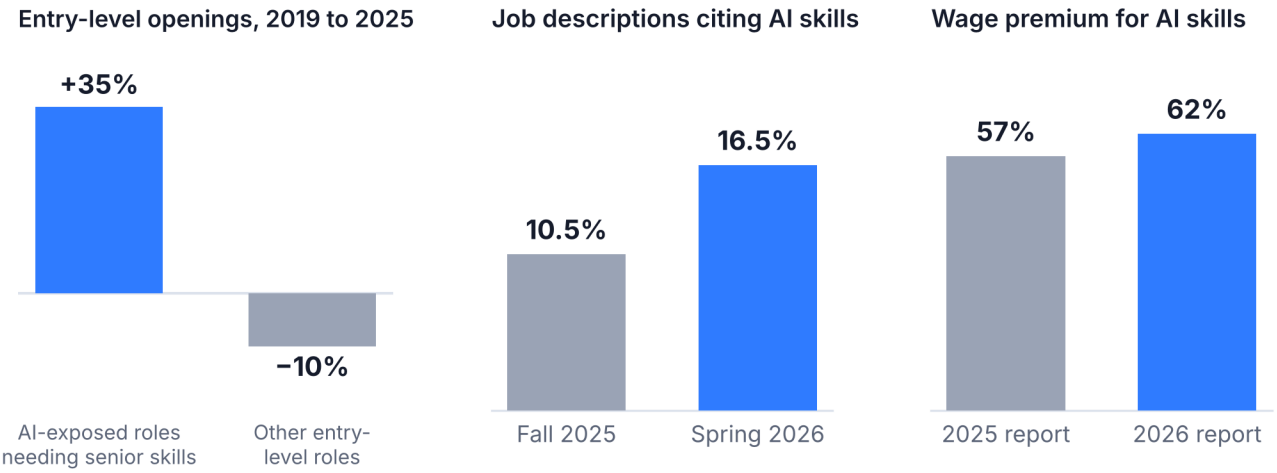


Chart 6. The 35% and 10% figures come from PwC's U.S. entry-level analysis. The wage premium is a global average. Source: [10] PwC, Jul 31, 2026 [12] NACE, Jul 22, 2026

SECTION 03

Value proposition

Each customer gets a different result. All three rest on the same product: measured practice that turns AI use into AI skill.

Customer	Job to be done	Gain we create	Proof we will show
Learner (15 to 22)	Get better results from AI for school, work, and side projects	Move from answer-seeking to drafting, checking, and building. Leave with a portfolio.	Rubric score gain, artifacts shipped, iteration depth
Program director or college lead	Deliver a credible AI skills program that funders will back	Ready-to-run cohort with facilitator guide and an outcomes report	Blind-scored pre and post results, retention, cost per learner
District leader	Meet AI literacy expectations with low teacher lift and clear evidence	Standards-mapped practice layer with admin reporting	Pilot evidence and exportable usage data, which buyers now ask for [23]

POSITIONING STATEMENT

For students ages 15 to 22 from communities schools underserve, who already use AI but were never taught to use it well, TheDirectLine.AI is the practice platform that builds measurable AI skill through real projects and graded feedback.
Unlike free AI literacy curricula [19] [20], which teach about AI, DirectLine has students do the work, scores it, and proves the gain.

Why the gap matters for our framing

Employers rank AI literacy last of seven skills for new graduates (3.6 of 5) and critical thinking first (4.3) [13]. We do not sell "AI literacy." We sell judgment, communication, and proof of work, built through AI. That framing matches what employers value and what the PwC data shows they pay for [10].

SECTION 04

Customer segmentation

In education, the user and the buyer are different people. We segment both and pick a beachhead where a buyer exists and the sales cycle is short.

Learner segments (users)

Segment	Ages	Context and constraints	Who pays	Role in plan
S1. High school students in Title I schools	15 to 18	School devices and content filters. Parental consent and FERPA apply. Some districts restrict student AI use.	District or school	Long-term core
S2. Students in community programs (college access, after-school, summer)	15 to 18	Program staff deliver. Flexible time. Guardian consent handled by the program.	Nonprofit or funder	Beachhead
S3. Community college and first-gen university students	18 to 22	Adults who consent for themselves. Work and family demands. Phone-heavy.	College, student support programs	Beachhead
S4. Young professionals and workforce trainees	18 to 22	Job outcomes drive enrollment. Short, intensive formats.	Workforce boards, employers	Expansion

Buyer and influencer segments

Buyer	What they own	What they need to say yes
B1. District leaders (CAO, CTO, CTE director)	Curriculum, tech budget, data privacy agreements	Compliance, outcome evidence, policy alignment, low teacher lift

Buyer	What they own	What they need to say yes
B2. Teachers and librarians	Classroom time, adoption	Ready lessons, under one hour of prep, visible student progress
B3. Nonprofit program directors	Program design, funder reporting	Measurable outcomes for funders, easy facilitation, cohort pricing
B4. Community college deans and support program leads	Student success programs	Retention and career outcomes, adult-appropriate content
B5. Funders (foundations, workforce boards, corporate CSR)	Grants and sponsorships	Reach into target communities, credible evaluation, cost per learner

BEACHHEAD RECOMMENDATION

Start with S2 and S3, sold to B3 and B4. These buyers move in weeks, not school years. S3 learners consent for themselves. Both segments sit outside the K-8 bans now in place in New York City and on LAUSD devices [15] [16]. Results from these cohorts become the evidence B1 district leaders require.

SECTION 05

Pain points

Each pain point is a hypothesis until interviews confirm it. Severity is our current estimate.

Stakeholder	Pain point	Evidence today	Severity
Learners	No clear starting point beyond asking a chatbot a question	Named as a barrier by FGLI students [6]	High
Learners	Cannot tell when AI output is wrong or generic	Lower conceptual literacy among Black and Hispanic undergrads [7]	High
Learners	School does not teach AI well enough	Only 36% of teens say school teaches enough [5]	High
Learners	Unclear rules and fear of cheating accusations	NYC and LAUSD restricted student AI use in Sep 2026 [15] [16]	Medium
Learners	No proof of AI skill for colleges or employers	1 in 3 entry-level jobs now requires AI skills [11]	Medium
Learners	Paid AI tiers and good devices cost money	Resource constraints named by FGLI students [6]	Medium
Teachers	No training and no time in the schedule	High-poverty districts train teachers at 39% vs 67% [8]. Scheduling is hard [17]	High
Program directors	Funders want measurable outcomes, staff lack AI expertise	Hypothesis. Test in interviews	High (untested)

Stakeholder	Pain point	Evidence today	Severity
District leaders	Unclear guidance, privacy review, proof of efficacy	52% of teachers report no or unclear guidance [9]. Contracts rarely define success [23]	High
Employers	Entry-level hires lack judgment with AI tools	AI-exposed entry roles 7x more likely to need senior skills [10]	Medium

SECTION 06

Assumptions and how we grade them

We grade every assumption on two axes, then test the riskiest first. Interviews exist to move these scores, not to collect compliments.

- **Importance (1 to 5).** If this is false, how badly does the plan break? 5 means the company fails.
- **Evidence (1 to 5).** How much proof do we have today? 1 means none. 5 means strong primary data.
- **Risk score** = Importance x (6 minus Evidence). Range 1 to 25. Test everything at 16 or above first.
- **Pass threshold.** Set before the test. If the result misses, we change the plan, not the threshold.

Willingness to pay and feedback accuracy carry the most risk

Risk score = Importance x (6 minus Evidence). Range 1 to 25.

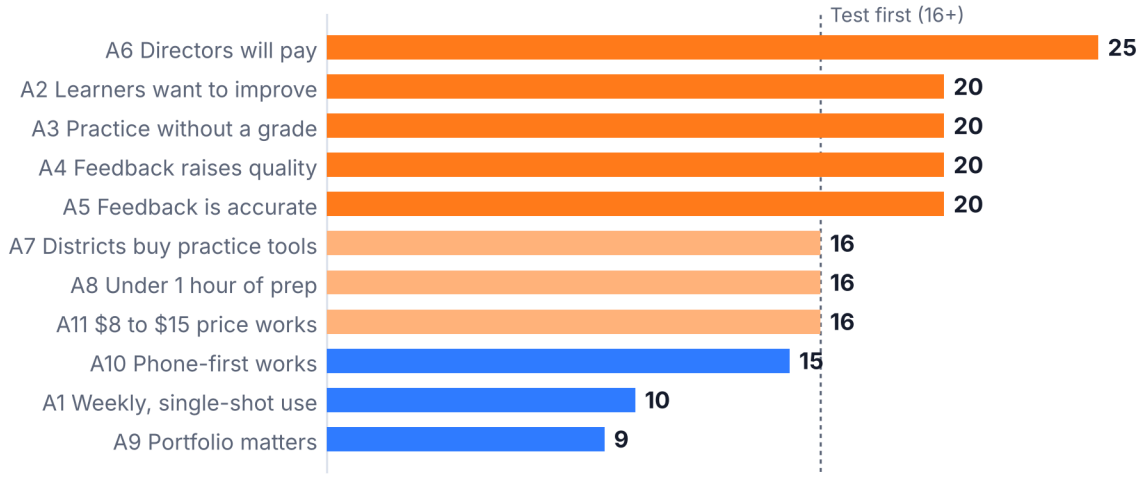


Chart 7. Scores reflect evidence available on the date of this version. A1 scores low risk because Pew, College Board, and Britebound already support it. Source: [1] Pew Research Center, Feb 24, 2026 [4] College Board, Dec 11, 2025 [5] Britebound, Jun 2026

#	Assumption	Imp	Evid	Risk	Test	Pass threshold
A6	Program directors will pay for or fund a cohort program	5	1	**25**	10 director interviews with a letter of intent ask	3 LOIs or 2 paid pilots
A2	Learners want to get better at AI, not just get answers faster	5	2	**20**	Learner interviews plus a waitlist smoke test	6 of 10 name a goal and commit 30 min a week

#	Assumption	Imp	Evid	Risk	Test	Pass threshold
A3	Learners will practice without a grade attached	4	1	**20**	Two-week concierge cohort	40% finish 3+ challenges unprompted
A4	Rubric feedback raises output quality measurably	5	2	**20**	Pre and post challenge, blind-scored by two educators	Mean rubric score up 25% after 3 iterations
A5	AI feedback can be accurate and free of fabrication	5	2	**20**	100-attempt eval set graded by educators and model	85% agreement within 1 point, zero fabricated facts
A7	Districts will buy a practice tool, not only curriculum	4	2	**16**	5 district leader interviews	3 of 5 name a budget line and an owner
A8	Staff can run a session with under one hour of prep	4	2	**16**	Pilot facilitator test	Prep under 1 hour, facilitator rating 4 of 5
A11	$8 to $15 per learner per year is an acceptable district price	4	2	**16**	Price sensitivity questions in buyer interviews	Acceptable range overlaps ours
A10	Learners can complete artifact challenges on a phone	3	1	**15**	Device survey and usability test	70% complete a challenge on mobile
A1	Target learners use AI weekly, mostly single-shot	5	4	**10**	Ask for the last 3 times they used AI	7 of 10 weekly, 6 of 10 single-shot
A9	Colleges and employers value an AI portfolio	3	3	**9**	5 recruiter and admissions interviews	3 of 5 say they would review it

Shading: dark orange is risk 20 or above, light orange is 16 to 19, blue is below 16.

WHAT THIS TELLS US

A6 decides whether we have a business. A5 decides whether we deserve one. Published studies show zero-shot AI grading agrees poorly with human raters [24] [25], so A5 is a real risk, not a formality. Both get tested in the first 60 days.

SECTION 07

Customer discovery plan

30 interviews in four weeks. Each interview maps to assumption IDs from Section 06. We ask about past behavior, not opinions about the future.

Group	Count	Recruit through	Assumptions tested
Learners 15 to 17 (via programs, guardian consent)	6	College-access nonprofits, MIT RAISE contacts	A1, A2, A3, A10
Learners 18 to 22	6	Community colleges, first-gen student groups, Consortium and Kappa Alpha Psi networks	A1, A2, A3, A9, A10
Teachers and librarians	6	Boston-area schools, Harvard i-lab network	A8, A4
Program directors and college success leads	6	Nonprofits, student support programs, REDF network	A6, A8, A11
District leaders	3	HKS and Tuck alumni, i-lab advisors	A7, A11
Recruiters and admissions staff	3	Tuck and HKS career networks	A9

Interview rules

- Ask about the last specific time, not what they usually do.
- No pitching until the last five minutes. Then show one screen and watch.
- Record with consent. For anyone under 18, get guardian consent and keep a program staff member present.
- Two people per interview when possible: one asks, one takes notes.
- Tag every note to an assumption ID. Re-score Section 06 every Friday.

Learner interview guide (30 minutes)

- Walk me through the last three times you used an AI tool. What were you trying to do?
- Show me the prompt you typed. What did you do with the answer?
- Tell me about a time the AI gave you something wrong or useless. How did you know?
- What would you build or make if you knew how to use these tools well?
- Who taught you what you know about AI? What did they get right or wrong?
- What rules does your school or job have about AI? How do you feel about them?
- What device and internet do you use for school or work most days?
- If a free program met weekly for six weeks, what would make you stay?

Program director and educator guide (30 minutes)

- What did your students or participants do with AI last month? How do you know?
- What have you tried to teach about AI so far? What happened?
- How do you report outcomes to funders today? Which metrics matter most?

- Walk me through how you bought the last tool or curriculum. Who signed, and how long did it take?
- What budget line would a program like this come from? What did you pay for the last comparable program?
- What would make you say no in the first meeting?
- Would you sign a letter of intent for a pilot next semester if the results met a bar you set?

District leader guide (30 minutes)

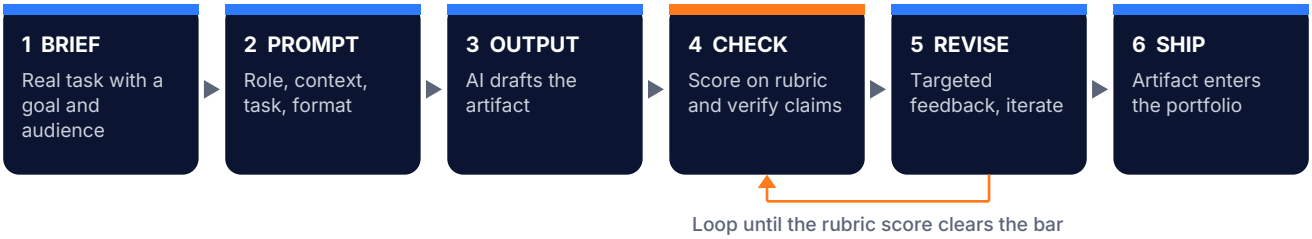
- What does your current AI policy require, and what does state guidance now ask of you?
- What AI tools are approved today? How did they pass privacy review?
- How do you measure whether an edtech tool worked? What evidence did your last renewal require?
- Who owns AI literacy in the district: curriculum, CTE, libraries, or technology?
- At what price per student would this feel cheap, fair, expensive, and too expensive?
- What would a pilot need to show for you to scale it?

SECTION 08

Product: what we build

We build the practice and feedback layer that free curricula leave out. Students learn by doing real work with real tools and getting graded on it.

The core loop



Where students are, and where we take them

Kim (2026) describes AI adoption as a depth gradient with four modes and finds FGLI students mostly stuck at the first [6]. Our skill ladder is built to move learners down this table.

Adoption mode (Kim, 2026)	What it looks like	DirectLine level
1. Basic chatbot interface	Ask a question, copy the answer. Where most FGLI students stop.	L1 Ask, L2 Check
2. Tool-augmented prebuilt interface	Files, search, voice, image, and artifact tools inside an AI app	L3 Iterate, L4 Build
3. Agentic development interface	Tools like Claude Code or Replit that build multi-step projects	L4 Build, L5 Direct
4. Programmatic integration	Calling models from code inside a workflow or product	L5 Direct (career track)

Design principles

- **Practice over lectures.** Every session ends with a finished artifact, not a quiz.
- **Judge before you trust.** Verification is a required step, not an optional tip.
- **Tool-agnostic skills.** Skills transfer across Claude, ChatGPT, Replit, ElevenLabs, and Wispr Flow. Challenges run in our interface, then repeat in the real tool.
- **Phone first.** Every challenge must work on a phone over a slow connection.
- **Safe by default.** No training on student data, content moderation, and an educator view of every session.

Skill ladder (curriculum spine)

Level	Skill	Learner can	Example artifact
L1 Ask	Structured prompting	Give role, context, task, constraints, and format in one request	A tailored cover letter draft
L2 Check	Verification	Spot errors, test claims against sources, flag weak reasoning	A fact-checked research brief
L3 Iterate	Refinement	Compare outputs, give targeted feedback, reach a quality bar	A revised personal statement with change log
L4 Build	Multi-step artifacts	Plan and produce a deck, app, audio piece, or data analysis	A 5-slide pitch deck or a simple web app
L5 Direct	Workflows across tools	Chain tools, reuse prompts, delegate tasks to agents, review the result	A small business workflow using voice, text, and code tools

Quality and anti-hallucination controls

Trust is the product. A tool that teaches students to verify AI cannot give them wrong feedback. Research shows why we cannot rely on an AI grader out of the box: in one benchmark, ChatGPT's scores correlated with human scores at r=0.23, while two human raters agreed at r=0.72 [24]. A second study of five models found human-AI agreement low and non-significant [25]. Prompting with grade level and examples helped in some cases [24]. That is the basis for our controls.

Zero-shot AI grading falls short of human agreement

Correlation with human essay scores, ASAP Task 1 (1,783 argumentative essays). Best reported values shown.

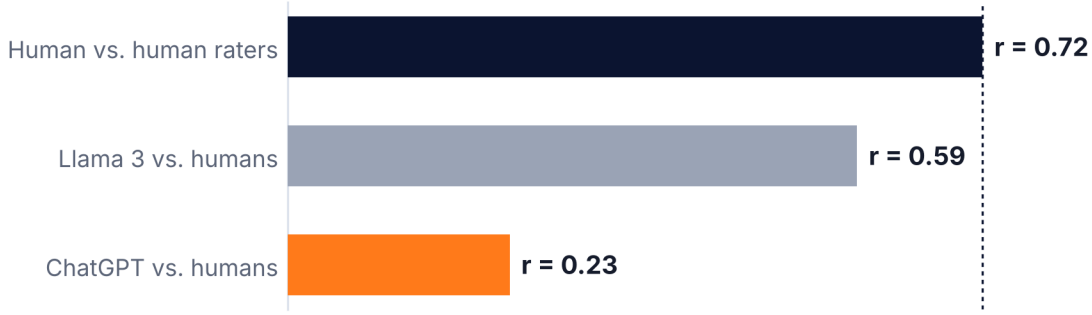


Chart 8. Zero-shot grading is the failure case our controls are designed to prevent. Source: [24] Kundu and Barbosa, arXiv 2409.13120, 2024

Control	How it works	Metric
Rubric-bound grading	The model scores only against a fixed rubric per challenge. It cannot introduce new factual claims in feedback.	Feedback stays on rubric criteria in 100% of audited samples
Educator exemplars	Educators write a strong and a weak reference answer for every challenge. The model grades against both.	Every challenge has 2+ exemplars before launch
Eval harness	Each challenge runs against 100+ educator-graded attempts before release and after every model change.	85%+ agreement within 1 rubric point
Citation check	Research challenges require sources. The platform checks that links resolve and flags quotes it cannot find.	Unverifiable citations flagged before scoring
Human audit	Educators review a random 10% of graded sessions each week, as both studies recommend [24] [25].	Error rate trend, reported monthly
Teach the failure	Some challenges contain planted AI errors. Learners score points for catching them.	Catch rate rises across L2 to L4

Curriculum alignment

Each challenge maps to the ISTE Standards for Students and the AI4K12 Five Big Ideas, plus state AI guidance where it exists. Ages 15 to 17 get a guarded track with school-safe topics. Ages 18 to 22 get a career track with job-search, workplace, and entrepreneurship artifacts. MIT Technology Review's call for prompt-based workflow skills, verification skills, and domain judgment maps directly to L1, L2, and L5 [14].

SECTION 09

Three approaches to the goal

All three share the core loop and quality controls. They differ in who we sell to first, how fast we learn, and how much evidence we produce.

APPROACH A

Studio: direct to learners

A free or freemium app learners join on their own. Growth comes from social content, referrals, and ThroughLine.AI channels.

STRENGTHS

- Fastest product learning and direct data
- Builds a public brand and a waitlist
- Low cost to start

WEAKNESSES

- Retention without a grade or cohort is unproven (A3)
- Minors need consent flows
- Students have little ability to pay

APPROACH B

Program-in-a-Box: licensed cohorts

A 6 to 8 week cohort program for nonprofits, college-access groups, and community colleges. Includes the platform, facilitator guide, challenges, and an outcomes report for funders.

STRENGTHS

- A buyer exists with a budget and funders
- Short sales cycles and consented learners
- Produces blind-scored outcome data

WEAKNESSES

- Services-heavy at first
- Depends on partner staff quality
- Smaller contracts than districts

APPROACH C

Classroom Layer: district platform

A teacher and admin product inside Google Classroom or Canvas. Standards-aligned units, a teacher dashboard, and district reports.

STRENGTHS

- Largest contracts and the long-term goal
- Fits district demand for AI literacy
- Reaches whole schools at once

WEAKNESSES

- 12 to 18 month sales cycles
- Heavy privacy and compliance review
- NYC and LAUSD restrict student AI use [15] [16]

Side-by-side comparison

Criterion	A. Studio	B. Program-in-a-Box	C. Classroom Layer
Time to first revenue	Slow	Fast (1 to 3 months)	Slow (12+ months)
Outcome evidence produced	Medium	High	High, but late
Build cost to start	Low	Medium	High
Compliance load	Medium	Low to medium	High
Scale ceiling	Medium	Medium	High
Fit with adoption goal	Medium	High	High
Tests the riskiest assumptions (A6, A5)	Partly	Yes	Partly

RECOMMENDATION

Lead with B. Build the A core loop as the engine. Target C for the 2027 to 2028 school year.

B gets us paying partners, consented learners, and blind-scored results within one semester. Day of AI's Maryland pilot shows the format districts accept: a defined cohort in high-need schools with outcome reporting [18]. Keep a light A tier free so pilot graduates keep practicing.

SECTION 10

Evidence exhibits

These exhibits test the core hypothesis. Each figure was checked against the linked source on Oct 25, 2026. We include the evidence against us. Mentors should push hardest on Exhibits 11 to 14.

EXHIBIT 1

Black and Hispanic teens use AI for schoolwork more than peers

[1] Pew Research Center, Feb 24, 2026 [2] Pew Research Center, Feb 24, 2026 | Verified Oct 25, 2026

In fall 2025, 60% of Black teens and 62% of Hispanic teens used chatbots for schoolwork help, versus 49% of White teens. 38% of Black teens rated chatbots extremely or very helpful, versus 22% of White teens.

So what: Use is high in our segment. Supports A1.

EXHIBIT 2

The pattern holds across 10,000+ high school responses

[4] College Board, Dec 11, 2025 | Verified Oct 25, 2026

Black and Hispanic high school students were the most frequent GenAI users at 65%, compared with 58% of Asian and 55% of White students. Data collected June 2024 to June 2025.

So what: A second large sample confirms Exhibit 1.

EXHIBIT 3

Lower-income teens rely on AI hardest

[1] Pew Research Center, Feb 24, 2026 | Verified Oct 25, 2026

20% of teens in households under $30,000 do all or most of their schoolwork with chatbots, versus 7% in households at $75,000 or more. 43% of lowest-income parents have talked with their teen about AI, versus 56% at the top.

So what: Heavy reliance with less guidance at home. Replaces the secondary analysis cited in v0.1.

EXHIBIT 4

Use is not readiness

[5] Britebound, Jun 2026 | Verified Oct 25, 2026

Only 24% of teens rank as highly AI-ready. The share is 32% for White teens and 15% for Black and Hispanic teens. Note the nuance: in this survey, frequent overall AI use was 61% for Black teens and 69% for White teens, so the use advantage in Exhibits 1 and 2 is specific to schoolwork.

So what: The readiness gap is large and consistent. This is the gap DirectLine closes.

EXHIBIT 5

Access is not adoption

[6] H. Kim, arXiv 2609.36129, Sep 28, 2026 | Verified Oct 25, 2026

From 61 interviews, the author finds first-generation, low-income college students use AI mostly through chatbots for narrow tasks. Barriers: low perceived value, underestimated self-efficacy, an unclear starting point, low peer exposure, and resource constraints.

So what: Our thesis, stated by independent research. Each barrier maps to a product feature.

EXHIBIT 6

Conceptual literacy gap in college

[7] AERA 2026 Annual Meeting, Apr 2026 | Verified Oct 25, 2026

A survey of 700+ undergraduates at a large public university found Black and Hispanic students report lower conceptual GenAI literacy than White students.

So what: Use without understanding. Supports the L2 Check level.

EXHIBIT 7

High-poverty districts train fewer teachers

[8] K-12 Dive, reporting a RAND survey, Apr 18, 2025 | Verified Oct 25, 2026

In fall 2024, 67% of low-poverty districts trained teachers on AI, versus 39% of high-poverty districts. The overall share rose from 23% to 48% in one year.

So what: The instruction gap tracks poverty. Supports the beachhead in under-resourced communities.

EXHIBIT 8

Entry-level work now demands judgment

[10] PwC, Jul 31, 2026 | Verified Oct 25, 2026

AI-exposed entry-level roles are seven times more likely to require senior skills. Those roles grew 35% since 2019, while other entry-level roles fell 10%. The average AI skills wage premium reached 62%.

So what: Frame outcomes as career readiness, not prompt tricks.

EXHIBIT 9

Employers now screen for AI skills

[11] NACE, Apr 20, 2026 [12] NACE, Jul 22, 2026 | Verified Oct 25, 2026

Employers say more than one-third of entry-level jobs require AI skills, nearly triple fall 2025. Yet 50.5% of 17,000+ college students are not building AI skills.

So what: A skills gap with a hiring cost attached.

EXHIBIT 10

Experts call for verification and workflow skills

[14] MIT Technology Review, May 26, 2026 | Verified Oct 25, 2026

The authors argue colleges and training programs should build prompt-based workflow skills, verification skills, and domain judgment into ordinary degrees.

So what: Validates the curriculum spine in Section 08.

EXHIBIT 11

Against us: employers rank AI literacy last

[13] Strada Education Foundation, May 19, 2026 | Verified Oct 25, 2026

Among ~1,500 talent leaders, AI literacy ranked last of seven skills for new graduates (3.6 of 5). Critical thinking ranked first (4.3). 41% said AI has reduced foundational tasks for entry-level workers.

So what: Position DirectLine as judgment and communication through AI, not AI literacy as a label.

EXHIBIT 12

Against us: large districts are restricting student AI

[15] Education Week, Sep 3, 2026 [16] Associated Press via KSAT, Sep 2, 2026 | Verified Oct 25, 2026

New York City placed a one-year moratorium on student-facing generative AI through 8th grade. LAUSD blocked student generative AI on district devices for this school year.

So what: Start with ages 15 to 22 and program partners. NYC is adding AI literacy classes for high schoolers, which is an opening.

EXHIBIT 13

Against us: free curriculum is plentiful

[19] Day of AI, Accessed Oct 2026 [20] aiEDU, Accessed Oct 2026 | Verified Oct 25, 2026

Day of AI offers free curriculum under a noncommercial license. aiEDU offers free, standards-aligned K-12 materials, including a 10-week high school course.

So what: Do not sell curriculum. Sell practice, feedback, and measured outcomes. Partner where possible.

EXHIBIT 14

Risk: AI grading is unreliable out of the box

[24] Kundu and Barbosa, arXiv 2409.13120, 2024 [25] Gaggioli et al., arXiv 2508.02442, Aug 4, 2025 | Verified Oct 25, 2026

Two studies found weak agreement between LLM essay scores and human scores. Both recommend human supervision.

So what: Our quality controls are a requirement, not a feature. See Section 08.

SECTION 11

Business model and metrics

Prices below are hypotheses to test in interviews, anchored to market benchmarks. We will not quote them to buyers until A6 and A11 are tested.

Offer	Buyer	Pricing hypothesis	Basis
Program-in-a-Box cohort license	Nonprofits, community colleges	Per-cohort license for up to 30 learners. Range set in interviews.	Test with price sensitivity questions (A11)
District seat license	School districts	$8 to $15 per learner per year	Edtech AI typically runs $5 to $15 per student per year [21]. Source cites no data, so treat as directional.
Teacher and staff training	Districts, programs	Add-on per facilitator	Institutional AI tools run $8 to $30 per teacher per month [22]
Sponsored cohorts	Foundations, employers	Funder covers learners in a named community	Bridge while district deals mature
Free Studio tier	Learners	Free, limited challenges	Alumni practice and referrals

Metrics

Metric	Definition	Pilot target
North star	Learners who ship 3+ verified artifacts	60% of enrolled learners
Skill gain	Change in blind-scored rubric score, pre to post	+25% mean
Depth shift	Share of learners using AI beyond answer-seeking, pre to post	Answer-seeking-only share cut in half
Iteration depth	Average revisions per artifact	3 or more
Verification rate	Share of research artifacts with checked claims	90%
Retention	Learners active in week 6 of 8	70%
Facilitator lift	Prep time per session	Under 1 hour
Feedback accuracy	Model and educator agreement within 1 point	85%+
Unit cost	Model API cost per learner per cohort	Track from day one

The depth-shift metric reuses Britebound's answer-seeking measure [5], so pilot results compare directly to a national baseline.

SECTION 12

90-day roadmap

Fall at HKS covers discovery and the MVP. The pilot runs at the start of the Tuck spring term, with results ready for mentors by March.

Phase	Dates	Work	Exit criteria
1. Discover	Oct 26 to Nov 22	Run 30 interviews. Re-score assumptions weekly. Shortlist 3 pilot partners. Draft consent and data policy.	A1, A2, A6 re-scored. 1 partner LOI.
2. Build	Nov 23 to Dec 20	Build the MVP with Claude Code: 5 challenges (L1 to L3), rubric engine, eval harness, portfolio export, educator view.	Eval agreement at 85%+ on all 5 challenges.
3. Prepare	Dec 21 to Jan 17	Train the facilitator. Recruit 30 to 60 learners. Run a pre-assessment. Fix mobile issues.	Learners enrolled, baseline scores captured.
4. Pilot	Jan to Feb 2027	Run the 6 to 8 week cohort from Tuck. Weekly audit of 10% of graded sessions.	Week 6 retention and score gain measured.
5. Report	March 2027	Publish the outcomes report. Bring it to i-lab, Media Lab, and 3 districts.	2 paid pilots or district meetings booked.

SECTION 13

Risks, mentor asks, and open decisions

Key risks

Risk	Likelihood	Mitigation
Student data and safety incident with minors	Medium	Start with 18+ and program-run cohorts. No training on student data. Moderation and educator view. Legal review before launch.
AI feedback is wrong and damages trust	Medium	Rubric-bound grading, eval harness, 10% weekly human audit. Evidence: [24] [25]
Learners stop after week 2	High	Cohort format, real artifacts with personal value, facilitator check-ins.
Free curriculum makes us look redundant	Medium	Partner with Day of AI and MIT RAISE [19]. Sell practice and measured outcomes, not content.
District restrictions spread	Medium	Focus on ages 15 to 22 and program partners first [15] [16].
Long district sales cycles drain runway	High	Fund early work through program licenses and sponsored cohorts.
Model costs grow with usage	Medium	Route drafting to lower-cost models. Track cost per learner from day one.

Risk	Likelihood	Mitigation
Founder time across Tuck, HKS, and recruiting	High	One partner and one pilot before any new feature. Weekly 3-hour build block.

What we ask of mentors

Mentor group	Ask
Harvard Innovation Labs	Intros to Boston college-access nonprofits and community colleges. Feedback on entity structure and student data policy. Pressure-test of the pricing hypotheses.
MIT Media Lab and MIT RAISE	Help designing the pilot evaluation so results are credible. Intro to Day of AI to explore a partnership rather than compete.
Tuck	Classmates and faculty to review the business model and run price sensitivity tests in the spring term.

Open decisions

- **Structure.** DirectLine is a ThroughLine.AI product. Decide whether it needs its own entity, and whether a nonprofit arm should hold grant funding.
- **Beachhead geography.** Boston first, given i-lab and Media Lab access, or Dallas and Fort Worth, given founder networks.
- **Age floor for the pilot.** 18+ only, or include 15 to 17 through a partner with consent processes in place.
- **Model provider.** Choose a primary grading model based on eval results, not brand.

SECTION 14

Sources and verification log

Every source below was opened and checked on Oct 25, 2026. Click a title to open it. The last column states what was checked and how much weight the source can carry.

STATUS KEY

Verified: figure confirmed on the linked primary page or report. **Verified (abstract)**: confirmed in the published abstract, full paper not reviewed. **Verified as stated**: the page says it, but the page cites no underlying data. **Secondary, attributed**: confirmed as reported by a secondary outlet that names its source. Check the primary source before using it in a pitch or grant. **Framework reference**: standards cited for alignment, not for a figure.

#	Source	Figures checked	Status
[1]	How Teens Use and View AI (full report PDF) Pew Research Center, Feb 24, 2026	n=1,458 teens, fielded Sep 25 to Oct 9, 2025. 64% use chatbots, 54% for schoolwork, 10% do all or most schoolwork with them. Race and income breakouts used in Charts 1 and 2.	Verified
[2]	Demographic differences in how teens use and view AI Pew Research Center, Feb 24, 2026	Black and Hispanic teens more likely to use chatbots for schoolwork and rate them helpful.	Verified
[3]	About a quarter of U.S. teens have used ChatGPT for schoolwork Pew Research Center, Jan 15, 2025	13% (2023) to 26% (2024). Black and Hispanic 31%, White 22%. No difference by income.	Verified
[4]	Variation in high school attitudes toward GenAI College Board, Dec 11, 2025	10,000+ responses. Frequent GenAI use: Black and Hispanic 65%, Asian 58%, White 55%.	Verified
[5]	Teens and the AI Future: Perspectives, Possibilities, and Preparation Britebound, Jun 2026	n=3,000 teens grades 7 to 12. High AI-Ready: White 32%, Black 15%, Hispanic 15%. Answer-seeking only: Black 29%. 36% say school teaches enough.	Verified
[6]	Accessible, but Not Adopted: Increasing LLM Adoption among FGLI College Students H. Kim, arXiv 2609.36129, Sep 28, 2026	61 interviews. Use is shallow and limited to chatbots. Five barriers. Four-mode depth gradient.	Verified (abstract)
[7]	Beyond the Digital Divide: Sociodemographic Gradients in College Students' GenAI Usage AERA 2026 Annual Meeting, Apr 2026	700+ undergraduates. Black and Hispanic students report lower conceptual GenAI literacy.	Verified (abstract)
[8]	Teacher AI training remains uneven despite uptick K-12 Dive, reporting a RAND survey, Apr 18, 2025	Districts training teachers on AI: 23% (fall 2023) to 48% (fall 2024). Low-poverty 67%, high-poverty 39%.	Verified
[9]	AI Literacy Belongs in Every K-12 Classroom Unite.AI, citing Ipsos and CDT, Oct 9, 2026	Ipsos: 78% of teachers want responsible AI use in curriculum, 52% report no or unclear school guidance.	Secondary, attributed
[10]	2026 AI Jobs Barometer press release PwC, Jul 31, 2026	AI skills wage premium 62% (57% prior year). AI-exposed entry-level roles 7x more likely to need senior skills. +35% vs -10% openings (U.S.).	Verified

#	Source	Figures checked	Status
[11]	Demand for AI skills in entry-level jobs nearly triples since fall 2025 NACE, Apr 20, 2026	Employers say more than one-third of entry-level jobs require AI skills. n=185 employers.	Verified
[12]	Ready or Reluctant? Employer Expectations for AI Skills Meet Student Skepticism NACE, Jul 22, 2026	AI skills in job descriptions 10.5% to 16.5%. 50.5% of 17,000+ college students are not building AI skills.	Verified
[13]	Entry-Level Hiring in the AI Era: What Employers Are Thinking (and Doing) Strada Education Foundation, May 19, 2026	~1,500 leaders. AI literacy ranks last of 7 skills (3.6 of 5). Critical thinking first (4.3). 41% say AI cut foundational tasks.	Verified
[14]	It's time to address the looming crisis in entry-level work MIT Technology Review, May 26, 2026	Calls for embedding prompt-based workflow skills, verification skills, and domain judgment in degrees.	Verified
[15]	LAUSD adopts moratorium on student AI use Education Week, Sep 3, 2026	Student generative AI blocked on district devices this school year.	Verified
[16]	NYC, the nation's largest school system, bans AI for students through 8th grade Associated Press via KSAT, Sep 2, 2026	One-year moratorium through grade 8. High schoolers get AI literacy classes and supervised pilots.	Verified
[17]	Take these 4 steps when rolling out AI literacy lessons Education Week, Nov 2025	Districts struggle to fit AI literacy into the schedule and face teacher pushback.	Verified (excerpt)
[18]	Maryland pilot announcement Day of AI, 2025	5,000+ K-5 students, 12 schools, 60%+ free or reduced-price meals, 61% Black, Hispanic, or multiracial.	Verified
[19]	About Us Day of AI, Accessed Oct 2026	Free curriculum under CC BY-NC-SA. Independent nonprofit since Aug 2024, started at MIT RAISE.	Verified
[20]	AI Literacy Curriculum for Every Classroom aiEDU, Accessed Oct 2026	Free, standards-aligned K-12 materials, including a 10-week high school course.	Verified
[21]	AI Cost Math for K-12 Districts: Per-Seat vs Usage-Based in 2026 ibl.ai, May 30, 2026	Typical edtech AI runs $5 to $15 per student per year. Author cites no underlying source.	Verified as stated
[22]	The Real Cost of AI Tools for Schools OpenEduCat, 2026	Institutional AI tools typically run $8 to $30 per teacher per month.	Verified as stated
[23]	Schools are buying AI fast, but their contracts still can't show it helps MarketScale, 2026	Many districts do not define success in AI contracts. Recommends exportable usage reporting.	Verified
[24]	Are Large Language Models Good Essay Graders? Kundu and Barbosa, arXiv 2409.13120, 2024	ASAP Task 1: human-human r=0.72, Llama r~0.59, ChatGPT r~0.23. LLMs need human supervision.	Verified
[25]	Reliability and Validity of LLMs for Automated Assessment of Student Essays Gaggioli et al., arXiv 2508.02442, Aug 4, 2025	Five models, 67 essays. Human-LLM agreement low and non-significant. Human oversight critical.	Verified
[26]	ISTE Standards for Students ISTE, Framework	Curriculum alignment reference.	Framework reference
[27]	Five Big Ideas in AI AI4K12, Framework	Curriculum alignment reference.	Framework reference

Removed in v0.2

- 68% of urban teachers with no AI training and the count of 27 states with AI education bills. The cited blog page did not show these figures when checked.
- More than 80% of students not taught AI by teachers. The article that reported it names no source. Replaced with Britebound's 36% figure [5].
- Maryland AI Ready Schools Act details. Not verified in this pass. Removed until checked against the statute.
- A third-party analysis of Pew income data. Replaced with figures taken directly from the Pew report [1].

APPENDIX

ThroughLine.AI deliverable standard

Every DirectLine document, deck, and one-pager follows this standard so the work reads as one ThroughLine.AI brand.

Color

Deep Space	Navy	Signal Blue	Horizon Orange	Chrome Silver	Ink
#050812	#0B1430	#2F7BFF	#FF7A1A	#C9D1DC	#151B2B

Element	Standard
Covers and title slides	Deep Space background with the TL chrome mark or brand image. White headline, Chrome Silver subhead, Horizon Orange eyebrow label.
Interior pages	White background. Deep Space header band with the TL mark. Blue to orange gradient rule under every section title.
Typeface	Inter. ExtraBold for titles, Bold for headings, Regular for body, SemiBold for table labels.
Charts	Title states the takeaway. Horizon Orange marks the communities we serve, Slate marks the comparison group, Signal Blue marks the DirectLine outcome or growth. Direct labels, no gridlines, sample size in the caption.
Emphasis boxes	Navy box for core claims and recommendations. Blue-tint box for guidance. Warm-tint box with orange bar for risks and counter-evidence.
Writing	ToughGong style: short sentences, active voice, direct address, claims backed by data. No em dashes, semicolons, or hashtags.
Citations	Every figure carries a clickable bracketed number. A verification log with status closes every research document.
Lockup	"A ThroughLine.AI product" appears on every cover and footer.