



AI-powered math learning — built for GED math, engineered to extend beyond it.

OVERVIEW

MathPath AI™ is an AI-powered math learning platform that combines a structured curriculum, independently verified AI-generated questions and tests, context-aware and conceptual AI tutoring, and educator management tools in a single integrated web application. It is built for GED math today and engineered to extend beyond it — the same content and verification machinery applies to the SAT (adding advanced algebra and introductory right-triangle trigonometry) and, in principle, any math curriculum or level.

The platform serves students, tutors, teachers, and platform administrators through role-specific experiences built on a shared, self-sustaining content infrastructure. It currently covers the complete GED Mathematical Reasoning subject area — **27 lessons** across **4 subject areas**, **169 mapped concepts**, and a three-tier question bank — and is production-ready and live at mathpathai.com. (A right-triangle trigonometry subject is in active development as the first step of SAT/ACT coverage.)

TWO PROPRIETARY CAPABILITIES

Reliable, conceptual AI math is the gating capability for any math product — and the hardest part to get right. MathPath AI is a working, validated implementation of two content-agnostic methods:

- **C²AI™ (Contextual, Conceptual AI)** — the tutoring model. It knows the lesson the student is on, the specific problem, and the answer they got wrong, and teaches the underlying concept and *why* the math works — not just the steps. Calculators are tools, not teachers.
- **IVQG™ (Independently Verified Question Generation)** — the content pipeline that lets a question bank sustain 99%+ audited accuracy with no manual content treadmill.

CONTENT ARCHITECTURE

The question bank is organized as an **81-slot matrix**: 27 subtopics across 3 difficulty tiers — Foundational, Test Ready, and Challenge. Each slot maintains a configurable capacity target with per-slot overrides for topics that have natural content ceilings. The live bank holds **5,000+ verified questions**.

All 169 concepts are systematically tagged across the bank. A dedicated concept index enables targeted practice (a student practicing “Slope-Intercept Form” is served only questions tagged to that concept), gap detection, and coverage reporting. Coverage is tracked continuously — the admin dashboard surfaces any concept that falls below minimum representation before a student ever encounters a gap.

CONTENT QUALITY ASSURANCE — THE IVQG™ PIPELINE

Every question is independently verified before it reaches the live bank, in two phases:

- **At generation** — each question is verified three ways by AI models from two companies (two separate Anthropic Claude model passes plus a Google Gemini model pass), all three of which must agree on the answer before the question is stored. Additional automated gates check the arithmetic, the explanation’s internal consistency, and answer-option integrity; anything that fails is dropped or routed to human review. No single model validates its own output.

- **Ongoing** — a weekly automated accuracy re-audit and a weekly similarity audit run against the existing bank, flagging any drift for review and archiving near-duplicate questions to preserve variety.

Independent accuracy: a full-bank blind-solve audit places live-bank accuracy at **99.25%**, with a 95% confidence interval of 97.7%–99.9% (Agresti–Coull). This is corroborated by additional measurements (98.5%, 98.7%, 99.6%) and the ongoing weekly check. The complete methodology and results are in the companion Accuracy Validation Report.

SELF-SUSTAINING CONTENT INFRASTRUCTURE

The bank operates autonomously with no routine manual maintenance. Several times a day a scheduled server process checks each slot against its capacity target and refresh interval: below-target slots have newly generated, verified questions added until they reach capacity, while at-target slots are refreshed — new verified questions replacing a matching number of the oldest — keeping content fresh for longer-term students. Because all generation and verification happen asynchronously, ahead of demand, AI-performance uncertainty is entirely absent from the student experience.

Administrators receive an email summary after every automated run, and a dashboard panel shows real-time bank health across all 81 slots. Human intervention is limited to approving flagged questions and adjusting per-slot capacity overrides.

MULTI-ROLE EDUCATOR INFRASTRUCTURE

- **Class management** — enrollment via a 6-character class code, whole-class and per-student lesson assignment, end-date tracking, co-teachers and assistants.
- **Teacher-constructed Tests** — build custom tests from the verified bank (filter by lesson, difficulty, question type, or keyword), publish (content frozen), and assign to classes or individual students with optional open/due dates and per-student scoring.
- **Per-student drill-down** — attempt history, per-question breakdown, missed-question review.
- **Usage and needs analytics** — active days, quizzes taken and passed, AI-tutor engagement, and the lessons where students struggle most.
- **Broadcast email** to enrolled students, and downloadable PowerPoint slideshows + vocabulary PDFs for all 27 lessons.

ENGINEERING QUALITY & DEPLOYMENT

- **269 automated API/unit tests** plus an end-to-end browser suite, all passing; zero build-time type errors.
- All AI calls are made server-side; no model credentials are accessible from the browser. Logged-in surfaces require a verified session.
- Continuous deployment to production on each approved change, with zero-downtime deploys.
- Automated monitoring confirms every scheduled process is running on time (UptimeRobot alerts on any lag); Rollbar captures and attributes all client- and server-side exceptions.

INTELLECTUAL PROPERTY NOTE

The specific methods, model-orchestration patterns, prompt architecture, and system design implementing the capabilities above are proprietary intellectual property of MathPath AI, LLC and are not disclosed in this brief. Neither C²AI™ nor IVQG™ is subject-specific — both are general AI-for-math IP.