

Fluency Intelligence

Making the invisible foundation of mathematics actionable

Students cannot give their full attention to fractions, ratios, expressions, and equations when basic calculation is still consuming scarce working memory. Fluency Intelligence is designed to help educators see that hidden constraint earlier and act on it responsibly.

Executive summary

District leaders face a familiar paradox: mathematics standards demand more reasoning, yet many students reach increasingly complex work while basic number relationships still require effortful reconstruction. In 2024, 39 percent of fourth graders and 28 percent of eighth graders performed at or above NAEP Proficient in mathematics. NAEP Proficient is a national performance standard, not a synonym for state-defined grade-level proficiency, but the trend still signals how much foundational work remains.¹

Fact fluency does not guarantee mathematical proficiency. It is one strand of a much larger braid that includes conceptual understanding, strategic competence, adaptive reasoning, and productive disposition. But when basic facts are accurate and efficiently available, students can devote more attention to the mathematics in front of them. When they are not, teachers often see the downstream symptoms — slow multi-step work, fragile application, avoidance, or anxiety — without a coherent view of the underlying fluency journey.^{2,6}

XtraMath Pro with Fluency Intelligence is designed to close that visibility-to-action gap. It builds on the short daily adaptive practice schools already use. Behind the teacher, the system assesses recent practice evidence, aligns supported patterns with research-informed strategies, audits every finding and recommendation for evidence and safety, and organizes the result into a concise weekly Instructional Brief. The teacher interprets the child and makes the decision.

By the end of middle school, fewer than **3 in 10** students are proficient in mathematics.

39% Grade 4

of fourth graders at or above NAEP Proficient

28% Grade 8

of eighth graders at or above NAEP Proficient

2024 NAEP Mathematics Results. Grade 4 and grade 8 are separate national samples in the same year, not one group of students followed over time. NAEP Proficient is a demanding national standard and is not a synonym for state-defined grade-level performance.

The district value proposition

- **See sooner.** Use frequent practice evidence to surface supported patterns before they compound.
- **Act more coherently.** Give teachers a prioritized brief rather than another dashboard to mine.
- **Preserve authority.** Keep AI beside the teacher, not between the teacher and the student.

1. Why fluency matters after elementary school

In this paper, fact fluency means accurate, efficient access to basic number relationships and the flexible use of what a student knows. It is not speed for its own sake, and it is not memorization detached from understanding.

Working memory is the bottleneck educators cannot see

Working memory is limited. Middle-school mathematics asks students to hold quantities, symbols, steps, and goals in mind while choosing strategies and checking whether an answer is reasonable. Cognitive-load research helps explain why performance can break down when too many processes compete for that limited capacity.^{3,4}

If a student must count, derive, or repeatedly reconstruct a basic fact, the calculation is not free. It uses attention that could otherwise support fraction equivalence, proportional reasoning, equation solving, or interpretation. Practice that develops automatic retrieval can reduce the demand of already-learned subtasks. The U.S. Department of Education's elementary mathematics practice guide describes automatic fact retrieval as one of many important goals and assigns strong evidence to carefully designed fluency-building activities embedded within broader instruction.²

The reinforcing risk: unfinished fact fluency consumes working memory → advanced work feels harder → anxiety and avoidance can rise → even less cognitive capacity is available.

A reinforcing risk, not a diagnosis: unfinished fluency can be one contributor among many.

The emotional cost can reinforce the cognitive cost

Math anxiety is not simply a reaction after a student gets an answer wrong. Research indicates that worry itself can consume working-memory resources, functioning like a second task running alongside the mathematics. Repeated overload can therefore make accessible ideas feel inaccessible and can erode a student's willingness to engage.⁵

A critical boundary

Slow work, avoidance, and weak performance have many possible causes. Practice data should never be treated as a diagnosis of disability, motivation, potential, home life, or complete mathematical understanding. The responsible opportunity is narrower: identify a supported fluency pattern, make uncertainty visible, and help the teacher decide whether and how to respond.

For a district, the benefit is not a new theory of achievement. It is earlier visibility into a consequential, addressable foundation — connected to the curriculum and interpreted by the educators who know the student.

2. From daily practice to instructional visibility

The signals are already being created. The scarce resource is the teacher time required to interpret them consistently.

An established practice habit

Since 2009, XtraMath has focused on a narrow but important job: helping students build math fact fluency through short, adaptive practice. In FY25, 6.2 million students spent 1.3 billion minutes practicing with XtraMath, according to the organization's annual report. That scale does not prove a specific learning effect. It does show that XtraMath is already a familiar part of the instructional routine in a large number of classrooms.⁷

That existing habit matters strategically. Students do not need another standalone testing event or a new conversation with a student-facing AI system. Each practice session can contribute a small piece of evidence about what appears secure, what is developing, and where performance is inconsistent. Across time, those signals can be more useful than a single snapshot — if they are interpreted with discipline.

Practice evidence is useful precisely because it is bounded

Daily practice can illuminate performance within the practice context: accuracy, consistency, change over time, and whether enough recent evidence exists to support a conclusion. It cannot reveal the whole child or replace curriculum-embedded assessment, diagnostic assessment, classroom observation, or professional judgment.

That distinction is a design advantage. A system that knows what it cannot infer can give educators a cleaner signal. Instead of asking a teacher to search rows of data for every student, Fluency Intelligence is intended to surface only supported patterns, show the evidence boundary, and focus attention where a human decision may be useful.

What becomes possible

- **Earlier visibility.** A developing pattern can be noticed before it appears only as failure in later work.
- **Less synthesis burden.** Teachers receive a prioritized interpretation rather than one more data display.
- **Instructional continuity.** A common evidence language can support classroom, intervention, and grade-level conversations.
- **No new student-facing AI experience.** Analysis happens behind the scenes; the teacher remains the instructional interface.

The product purpose


XtraMath builds the practice habit. Fluency Intelligence makes the evidence more usable. XtraMath Pro is

the access layer that brings that capability to classrooms and schools. Together, they turn a familiar student routine into a weekly cycle of instructional clarity without asking educators to surrender judgment or students to surrender another block of time.

This is a visibility layer, not a replacement for curriculum, intervention, assessment systems, or the professional knowledge of the teacher.

3. What Fluency Intelligence does

Fluency Intelligence is designed as a teacher-facing loop: **Practice → Assess → Align → Audit → Brief**. Its value comes from the whole sequence. Analysis without alignment creates another data problem. Recommendations without audit create a trust problem. A brief without teacher authority creates an instructional problem.



The learning loop. Ten minutes of daily student practice becomes ten minutes of weekly teacher reading; the analysis, alignment, and auditing stay hidden, server-side. Teacher instruction closes the loop.

Practice	Students continue the short, familiar XtraMath routine. No chatbot interrupts practice; no student explains personal context to an AI system.
Assess	Recent practice evidence is ingested, aggregated, and classified into supported patterns — starting with a question many analytics products skip: is there enough evidence to say anything useful?
Align	Supported needs are compiled against research-informed mathematical strategies and instructional considerations — to narrow the teacher's next decision, not to generate generic advice.
Audit	Every finding and recommendation is evaluated for unsupported claims, contradictions, invented information, deficit framing, diagnostic language, and conclusions that exceed the evidence — and the result is logged. A failed output is withheld or replaced, not merely softened.
Brief	The teacher receives a weekly Instructional Brief: the patterns worth attention, the evidence behind them, relevant strategy options, and the limits of what can be concluded.

The Instructional Brief should help a teacher decide where to look, whom to confer with, and what to try — not prescribe a verdict. The system performs the synthesis. The teacher interprets the child and decides.

Inside the hidden stages

The three server-side stages each perform three bounded steps. None of them is visible to students or

teachers; together they are the work that turns raw practice into a brief worth trusting.

Assess Pattern analysis

The system reads the week's practice evidence and asks, before anything else, whether there is enough of it to support a conclusion.

- | | |
|------------------|--|
| INGEST | Collects the week's practice attempts — correct and incorrect answers, response times, which facts were practiced, and when. |
| AGGREGATE | Groups attempts by operation and fact family to reveal the shape of the week, rather than a scatter of individual data points. |
| CLASSIFY | Assigns each student only to states the evidence supports — fluent on some fact families, developing on others. |

Guardrails: Fluency Intelligence labels only what is in the data, and thin evidence is acknowledged rather than speculated over. The subject is always the pattern in the practice — never a judgment about the child.

Align Material alignment

Recommendations point to materials the school already owns, not to new content the teacher must adopt.

- | | |
|-----------------|--|
| COMPILE | Matches supported patterns to teaching strategies — accuracy practice, fluency-building activities, small-group discussion, and celebration moves. |
| RENDER | Translates those strategies into the school's own standards framework — CCSS, TEKS, or other state codes. |
| ASSEMBLE | Links each recommendation to the adopted curriculum — Illustrative Mathematics, Eureka, Bridges — so the next step is a page, not a purchase. |

Audit Governance audit

Safety checks run before anything reaches a teacher. A failed check blocks the output entirely — the language is not merely softened.

- | | |
|-----------------|---|
| PROCESS | Tags every claim in a draft finding to the specific evidence that supports it. |
| EVALUATE | Blocks diagnostic labels, judgments of teaching, deficit framing, and interpretations that exceed the evidence. |
| LOG | Records every audit decision and the full evidence chain, so any recommendation can be traced back to its source. |

What this looks like with real evidence

The excerpts below come from the 2026 beta, captured from a real classroom cascade. Names are pseudonymized; numbers and timings are preserved.

THE EVIDENCE · ONE STUDENT'S RECENT PRACTICE

Jordan · Addition · 251 attempts, 74% overall accuracy. The cumulative evidence flags three specific fact families:

- **Sums in the 6–10 range** (3+6, 2+4, 2+7): 56% accuracy at ~3.8 seconds per fact
- **Bridge 10 facts** (4+8, 3+8, 4+7): 62% accuracy at ~4.5 seconds
- **Plus One facts** (1+2, 1+9): 67% accuracy, faster at ~2.4 seconds

WHAT THE BRIEF CONCLUDES

"There's a prerequisite gap blocking further progress."

The report cites the three fact families by name, notes that "these patterns suggest the student may not have a reliable strategy for solving harder addition facts yet," and connects the gap to the specific standard involved — while also naming what is working: most facts answered correctly, and quickly.

No diagnosis. No character judgment. A supported pattern, its evidence, and a place for the teacher to look next.

At the classroom level, the brief organizes students into small strategy groups and aligns each group with a concrete, teacher-approved activity. From one week of the beta:

ADDITION

Making Ten

Building Toward Fluency
(Illustrative Mathematics)

David Baptiste · Bridge 10
Angela Rosales · Bridge 10
Areli Navarro-Hernan · Make 10
Javon Bell · Bridge 10
Fernando Martinez · Doubles

SUBTRACTION

Think Addition

Hitting the Target Number
(Illustrative Mathematics)

Aubree Rhodes · Doubles Reverse
Emily Stephen · Teen Bridge

DIVISION

Fact Family Triangles

Unknowns with Multiplication & Division (Khan Academy)

Kingsley Adebawale · Equal Groups

Strategy groups from a real beta Instructional Brief. Student names pseudonymized; groupings, strategy classifications, and aligned activities as generated.

The practical benefit

For teachers, the gain is instructional clarity without another dashboard workflow. For school leaders, it is a more consistent evidence-to-action process. For districts, it is a responsible way to make an existing investment in daily practice more useful while keeping governance visible and professional authority intact.

4. A responsible district decision

The right question is not, "Does this use AI?" It is, "Does this give teachers better evidence, clearer choices, and stronger safeguards than they have today?"

Teacher authority and governance are product requirements

District trust cannot depend on a promise that the system will usually be careful. The product must make its refusals as visible as its recommendations. At minimum, Fluency Intelligence should:

- **Stay inside the evidence.** No diagnosis, clinical label, character judgment, or invented context.
- **Show uncertainty.** Insufficient or conflicting evidence must be named, not hidden.
- **Block unsafe output.** Unsupported recommendations and contradictions should be withheld or replaced.
- **Keep the teacher in charge.** The brief informs professional judgment; it does not automate placement or intervention.

Separate what is established from what must be learned

The learning-science foundation is credible: working memory is limited; automatic retrieval can reduce the demand of already-learned subtasks; anxiety can consume cognitive resources; and fluency should be accurate, efficient, flexible, and connected to understanding. What is not yet established is the impact of Fluency Intelligence itself on student achievement, teacher workload, intervention quality, or anxiety. Those are product hypotheses to test, not outcomes to advertise.^{2,3,4,5,6}

Five questions for a district pilot

1. **Instructional usefulness:** Do teachers find the brief accurate, timely, and actionable?
2. **Evidence traceability:** Can educators see why a pattern was surfaced and when evidence is insufficient?
3. **Teacher agency:** Does the brief improve judgment without becoming a prescription or compliance burden?
4. **Equity and safety:** Are false positives, missing data, and subgroup differences monitored and addressed?
5. **Measurable value:** What changes in time-to-insight, teacher action, fluency development, and later readiness?

Availability and procurement

As of October 2026, Fluency Intelligence is available through XtraMath Pro for the 2026–27 school year at ~\$2 per student per year — see the [XtraMath pricing page](#) for current plans. Districts should confirm

pricing and entitlements, plus accessibility, privacy/FERPA, data processing, rostering, professional learning, and support.

The case is simple: the system can do the repetitive synthesis. The teacher should keep the relationship, the context, and the decision.

Selected references

1. National Center for Education Statistics. *2024 NAEP Mathematics Results, Grades 4 and 8*. nationsreportcard.gov
2. Fuchs, L. S., et al. (2021). *Assisting Students Struggling with Mathematics: Intervention in the Elementary Grades*. IES/What Works Clearinghouse. ies.ed.gov
3. Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, 24(1), 87–185.
4. Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. *Cognitive Science*, 12(2), 257–285.
5. Ashcraft, M. H., & Krause, J. A. (2007). Working memory, math performance, and math anxiety. *Psychonomic Bulletin & Review*, 14(2), 243–248.
6. National Council of Teachers of Mathematics. (2023). *Procedural Fluency in Mathematics*. nctm.org
7. XtraMath. (2025). *FY25 Annual Report*.