

Vectored Instruction

A Breakthrough Technology in Modern Education

The Method That Helps Students Catch Up While Keeping Them on Grade Level

Dr. Roger E. Billings

President, International Academy of Science | Creator of Acellus

Journal of Science | journal.science.edu

August 2026

Vectored Instruction® is a registered trademark of Acellus / International Academy of Science.

Contents

Abstract

1. Introduction: The Grade-Level Dilemma

2. Background: The Problem Vectored Instruction Was Built to Solve

- 2.1 Knowledge is cumulative, and gaps compound
- 2.2 The classroom cannot go back
- 2.3 One-to-one tutoring works but does not scale
- 2.4 "Patching and hoping"

3. What Vectored Instruction Is and How It Works

- 3.1 The normal learning cycle
- 3.2 Detection: Deficiency Diagnostics
- 3.3 Repair: Customized Personal Instruction
- 3.4 Return: back to the current lesson, still on grade level
- 3.5 Visibility: the teacher console
- 3.6 Automatic in every deployment; continuously improving
- 3.7 Where Vectored Instruction sits in the Acellus system

4. Outcomes Reported by Districts Using Vectored Instruction

- 4.1 Le Grand Union Elementary School District, Merced County, California
- 4.2 Community School of Holy Angels (special education)
- 4.3 Lansing Unified School District 469, Kansas (Kansas Heartland Virtual Academy)
- 4.4 Enrollment and program outcomes
- 4.5 What these outcomes have in common

5. The Science of Why Vectored Instruction Works

- 5.1 Learning hierarchies and the primacy of prior knowledge
- 5.2 Mastery learning and the 2 sigma challenge
- 5.3 Cognitive load, schemas, and the zone of proximal development
- 5.4 Diagnostic assessment through the analysis of wrong answers
- 5.5 Formative assessment and the power of immediate, task-focused feedback
- 5.6 Retrieval practice and spacing
- 5.7 What the intelligent tutoring literature tells us
- 5.8 Acceleration, not remediation; grade level, not retention
- 5.9 Solving the scale problem
- 5.10 Motivation, self-efficacy, and the message the student hears

6. Why Vectored Instruction Is Especially Crucial for Students Studying at Home

7. Using Vectored Instruction in Your Teaching

- 7.1 For teachers and administrators using Acellus
- 7.2 Applying the principles without the platform

8. Applying Vectored Instruction Principles in Textbooks and Learning Resources

- 8.1 Build and publish the prerequisite map
- 8.2 Design assessments that diagnose, not merely score
- 8.3 Author bridge lessons, not review chapters
- 8.4 Keep the grade-level spine intact

8.5 Instrument the resource and improve it

9. Enhancing Vectored Instruction Going Forward

9.1 Sharper diagnosis

9.2 Richer instruction

9.3 Durable retention

9.4 A stronger teacher partnership

9.5 Evidence and continuous improvement

10. Conclusion

About the Author

References

Abstract

Every school district in the United States is bound by the same expectation: students must be taught the standards of the grade in which they are enrolled. Yet a large share of students arrive in each grade missing foundational knowledge that should have been mastered one, two, or even several years earlier. Because district policy rightly resists moving students backward, teachers have historically been left to "patch and hope," re-teaching fragments of earlier grades while attempting to cover the current curriculum. **Vectored Instruction®**, introduced by Acellus roughly a decade ago and now embedded automatically in every Acellus deployment, resolves this dilemma. When a student fails to demonstrate mastery of a grade-level lesson, a technology called **Deficiency Diagnostics** analyzes the pattern of the student's responses to identify the specific missing prerequisite. A second component, **Customized Personal Instruction**, then reaches back as many grade levels as necessary to rebuild that foundation before returning the student to the current lesson, all while the student remains enrolled in the grade-level course and on track for credit and graduation.

This paper describes how Vectored Instruction works, summarizes the outcomes districts are reporting, and, for the first time, situates the method within the educational and cognitive science literature that explains why it works: learning hierarchies and the cumulative structure of knowledge, mastery learning, cognitive load theory, diagnostic assessment through distractor analysis, formative feedback, and the intelligent tutoring literature. It then offers practical guidance for teachers, curriculum designers, and textbook authors who wish to apply the principles of Vectored Instruction, explains why the method is particularly indispensable for students studying independently at home, and identifies research-based directions for the next generation of the technology. A closing section introduces the CyberSchool Initiative, through which local districts can sponsor supported home-study programs built on Acellus and Vectored Instruction.

Keywords: Vectored Instruction, Deficiency Diagnostics, prerequisite knowledge, mastery learning, adaptive learning, learning gaps, grade-level instruction, independent study, CyberSchool

1. Introduction: The Grade-Level Dilemma

American education has spent decades debating how to help students who have fallen behind. The debate has produced summits, legislation, white papers, and committees. It has not produced a solution that operates inside the ordinary classroom, in real time, for every child. The 2024 National Assessment of Educational Progress (NAEP) reported that only 39 percent of fourth graders and about 28 percent of eighth graders performed at or above the Proficient level in mathematics, that a record one-third of eighth graders scored below Basic in reading, and that only 22 percent of twelfth graders were proficient in mathematics, with a record 45 percent below Basic (National Assessment Governing Board, 2025a, 2025b). Independent recovery trackers report that the average student still needs the equivalent of four to five additional months of schooling to reach pre-pandemic achievement, and that gaps between high- and low-poverty schools have widened, not narrowed (Lewis & Kuhfeld, 2024; Fahle, Kane, Reardon, & Staiger, 2024).

Behind these aggregate figures is a structural problem that every superintendent recognizes. A sixth-grade teacher is required to teach sixth-grade standards. But if a third of the class never mastered fourth-grade fractions, the teacher must choose between two unsatisfactory options: teach the grade-level material to students who cannot access it, or spend instructional time re-teaching earlier grades and never reach the current curriculum. Superintendent Scott Borba of Le Grand Union Elementary School District in Merced County, California, described the situation with unusual candor: "My sixth-grade teachers were trying to teach sixth-grade curriculum, but they couldn't. They were too busy going back and filling in holes in third, fourth, and fifth grade" (Acellus Journal, 2026a, p. 10). The result was, in his words, "a domino effect of academic regression cascading through every grade level."

The obvious remedy, placing the student back into the grade where the knowledge was missed, is not available. District policies, state law, and a substantial research base all discourage grade retention: the meta-analytic evidence shows no academic or socio-emotional benefit for retained students relative to promoted peers (Jimerson, 2001), and Hattie's synthesis places retention among the few common practices with a negative average effect on achievement (Hattie, 2009). Nor is the alternative, generic below-grade remediation, effective. Large-scale analyses show that students who are pulled into below-grade-level work spend hundreds of hours per year on assignments that do not advance them, and that students who instead receive grade-level content with just-in-time support on prerequisites complete substantially more grade-level work and struggle less (TNTP, 2018; TNTP & Zearn, 2021).

The dilemma, then, is this: the student must remain on grade level, but the student cannot learn grade-level material without foundational knowledge that lives in an earlier grade. Vectored Instruction was designed to dissolve this dilemma rather than to manage it. It keeps the student in the grade-level course, detects precisely which prerequisite is missing, delivers exactly the instruction needed to repair it, and returns the student to the lesson, typically within minutes and often without the student realizing that anything unusual has occurred.

This paper has four purposes. First, it describes Vectored Instruction in enough detail that a district leader unfamiliar with Acellus can understand how it works and can appreciate that it is not an add-on but an automatic component of every Acellus deployment. Second, it summarizes the outcomes that districts using the technology are reporting. Third, and most importantly for the research community, it examines the scientific literature that explains why the method works, drawing on more than fifty years of research on prerequisite knowledge, mastery learning, cognitive load, diagnostic assessment, feedback, and intelligent tutoring. Fourth, it offers guidance for practitioners and developers who wish to apply these principles in their own teaching, textbooks, and learning resources, and it identifies research-based directions for enhancing the technology going forward.

2. Background: The Problem Vectored Instruction Was Built to Solve

2.1 Knowledge is cumulative, and gaps compound

Mathematics, reading, science, and the study of history are cumulative disciplines: each new concept depends on concepts learned earlier. Gagné (1968) formalized this as a *learning hierarchy*, a structure in

which a target intellectual skill decomposes into ordered subordinate skills, each of which must be mastered before the skill above it can be acquired. When a student fails at the top of the hierarchy, the most likely explanation is not a defect in the current lesson or the student's ability, but a missing element lower in the structure. Ausubel (1968) stated the principle in a single sentence that has become one of the most quoted in educational psychology: "The most important single factor influencing learning is what the learner already knows. Ascertain this and teach him accordingly."

The empirical evidence for cumulative dependency is strong and specific. In two national longitudinal samples, Siegler and colleagues (2012) found that elementary students' knowledge of fractions and whole-number division uniquely predicted their algebra and overall mathematics achievement five to six years later, even after controlling for IQ, working memory, family income and education, and other mathematical knowledge. A fifth-grade fraction gap does not stay in fifth grade; it resurfaces as an algebra failure in tenth grade. Curriculum researchers have likewise shown that the highest-achieving nations organize their standards *coherently*, sequencing topics according to the logical structure of the discipline, whereas U.S. standards have historically been "a mile wide and an inch deep," which makes it easier for students to move forward with prerequisites unmastered (Schmidt, Wang, & McKnight, 2005).

2.2 The classroom cannot go back

A single teacher facing thirty students cannot pause the class to re-teach a fourth-grade concept to three students, especially when the remaining twenty-seven need something different. In a traditional setting, as the Acellus classroom guide observes, "either the student falls further behind, or the teacher must set up one-on-one tutoring to catch the student up with the rest of the class" (International Academy of Science, 2025, p. 20). Even the most committed teacher can seldom locate the precise missing concept for an individual student; doing so would require analyzing that student's pattern of errors, mapping the errors to a prerequisite structure that spans several grades, selecting the right instruction, and delivering it, all while teaching everyone else. Studies of teachers' ability to anticipate students' common wrong answers show that even experienced teachers predict them inconsistently, and that remediation authored from expert guesses rather than empirical error data does not reliably improve mastery (Gurung et al., 2023).

2.3 One-to-one tutoring works but does not scale

Bloom (1984) demonstrated that students tutored one-to-one under mastery conditions performed about two standard deviations above students in conventional classrooms, and he challenged the field to find methods of group instruction as effective as tutoring, the "2 sigma problem." Four decades later, meta-analyses confirm that high-dosage tutoring produces meaningful gains (pooled effect of roughly 0.29 to 0.37 standard deviations across randomized trials; Nickow, Oreopoulos, & Quan, 2024), but the same literature documents that effects shrink dramatically as programs scale: tutoring programs serving more than a thousand students average only about 0.16 standard deviations, because supply of skilled tutors, cost, and scheduling all bind (Kraft, Schueler, & Falken, 2024). With more than 425,000 teaching positions nationally either vacant or filled by under-qualified staff (Learning Policy Institute, 2026), the human capacity to diagnose and repair individual learning gaps at scale simply does not exist.

2.4 "Patching and hoping"

For these reasons, most schools have defaulted to what may be called patching and hoping: re-teaching the current lesson more slowly, louder, or with different words; assigning generic review packets; labeling the student as "behind" or "unmotivated"; and moving on because the pacing guide requires it. As one keynote address put it, "the student didn't fail because the lesson was bad. The student didn't fail because they aren't smart enough. The student failed because somewhere, maybe a year ago, maybe three years ago, they missed something. A foundational concept. A building block. And without that piece, nothing built on top of it will ever make sense" (Billings, 2026). Vectored Instruction was engineered to find that piece and to fix it, in the moment, for every student, without removing any student from grade level.

3. What Vectored Instruction Is and How It Works

Vectored Instruction is a closed-loop instructional technology built into the Acellus Learning System. The name is deliberate. A vector has both *magnitude* and *direction*. When a student is stuck, Vectored Instruction determines the direction of the intervention (which specific prerequisite concept, in which earlier course or grade) and its magnitude (how far back to go and how much instruction is required) and then applies exactly that intervention before returning the student to the point of difficulty. It is not remediation in the conventional sense of a separate course, a pull-out group, or a review packet. It is targeted, just-in-time reconstruction of foundational knowledge, delivered inside the grade-level course.

3.1 The normal learning cycle

In an Acellus course, a student is presented with digital lessons: professionally produced videos in which an expert teacher introduces one concept at a time. Each video is followed by assessments designed to confirm mastery of that concept before the student advances. Assessments range from comprehension and mastery items aligned to the lesson's objectives to instructive assessments that teach through guided discovery (International Academy of Science, 2025; Billings, 2025). Under ordinary conditions the student watches, demonstrates mastery, earns credit toward the course, and moves to the next concept. Progress is visible to the teacher in real time.

3.2 Detection: Deficiency Diagnostics

Vectored Instruction engages when the normal cycle breaks down, that is, when a student is unable to progress through the assessments for a lesson. The first stage is a powerful and complex diagnostic technology called Deficiency Diagnostics. Its purpose is not merely to register that an answer is wrong but to determine *why* it is wrong and, specifically, whether the failure stems from missing foundational knowledge that should have been acquired in an earlier grade.

Deficiency Diagnostics rests on an empirical discovery made over years of analyzing student responses across millions of learners: when students miss a question, they do not miss it randomly. Large percentages give the *same* wrong answers, and those patterned errors reveal the specific underlying concept the student is missing (Billings, 2026). Acellus therefore treats every incorrect response as data. Item distractors are engineered so that particular wrong choices correspond to particular misconceptions

or prerequisite deficits, and the system continuously gathers response data to determine why students are selecting a given wrong answer and where the specific learning gap lies (International Academy of Science, 2025, pp. 30–31). A related capability, Prism Diagnostics, takes a group of students who all missed the same item and separates them, as a prism separates white light, into distinct groups according to the different underlying deficiency each group has, so that "same wrong answer on the surface" leads to "completely different reasons underneath," and each student is routed accordingly (Billings, 2026).

Two levels of response are available. For a discrete misconception within the current lesson, the system can trigger a targeted help video keyed to the specific wrong answer the student chose. Notably, that video does not reveal the correct answer; it redirects the student's thinking and returns them to the item (International Academy of Science, 2025, p. 30). When the pattern of errors instead indicates that the student lacks a *prerequisite* concept, one that belongs to an earlier lesson, course, or grade, the system escalates to full Vectored Instruction.

3.3 Repair: Customized Personal Instruction

Once a deficiency has been identified, Acellus deploys Customized Personal Instruction: a sequence of instruction, drawn from the prerequisite curriculum, that reconstructs the missing foundation. This sequence may reach back one lesson or several grade levels. If a tenth grader is failing an algebra lesson because of a concept missed in fifth-grade fractions, Vectored Instruction reaches back to fifth grade, builds a targeted bridge of instruction to close that gap, and then brings the student forward again (Billings, 2026). Some students need a single step; others need several. Each intervention is delivered "at the precise intensity needed to repair the learning deficit" (International Academy of Science, 2025, p. 66). The prerequisite lessons use the same video-plus-assessment format as the rest of the course, so mastery of the prerequisite is confirmed, not assumed, before the student is returned.

3.4 Return: back to the current lesson, still on grade level

When the prerequisite has been mastered, the student is returned to the original lesson and completes the assessment again. Throughout, the student remains enrolled in the grade-level course, continues to earn credit, and remains on track for graduation. The entire loop typically occurs quickly, and, because the intervening lessons are blended into the flow of the course, the student usually does not realize that instruction not ordinarily part of the course has been supplied. This matters for more than convenience: it eliminates the labeling and loss of place that accompany traditional remediation, and it sends the student a very different message: "we know exactly what you're missing, we're going to fix it right now, and you're going to keep moving forward" (Billings, 2026).

A worked example from a live classroom

A student working through a mathematics course completes Step 19 (Properties) successfully but fails Steps 20 and 21 (Adding and Subtracting). Deficiency Diagnostics recognizes from the response pattern that the difficulty is not with the current procedure but with background material from a lower-grade course. Vectored Instruction presents the needed prerequisite lesson from that earlier course. After the intervention, the student returns to Step 21 and passes it with a perfect score (International Academy of Science, 2025, Figure 2.4).

In a U.S. Government course, a student repeatedly misses items on freedom of speech. The pattern indicates that the student does not understand why the Bill of Rights was created. Vectored Instruction inserts a lesson on the Bill of Rights from a prerequisite curriculum; once mastered, the student is returned to the freedom-of-speech lesson and completes the activities successfully (International Academy of Science, 2025, p. 31).

3.5 Visibility: the teacher console

Vectored Instruction operates automatically, but it does not operate invisibly to educators. From the Acellus teacher console (the Active Teacher Interface and Live Monitor), teachers and administrators can see exactly when a student encounters a need for Vectored Instruction, how the student became stuck, which prerequisite concept was diagnosed, which lesson was used to bridge the gap, and the student's "before" and "after" scores. A dedicated Vectored Instruction Intervention panel displays the total number of interventions the system has made in a class, the average percentage increase in scores following intervention, and recent examples; drilling down reveals per-concept and per-student detail (International Academy of Science, 2025, pp. 31–32, Figures 2.11–2.12 and 3.7). Teachers therefore watch the student get stuck, watch Vectored Instruction engage, watch the personalized instruction deliver, and watch the student charge forward again, completing the lesson successfully. The dashboard also sorts students by who needs help most, so the teacher's scarce attention is directed where the software cannot substitute for a human.

3.6 Automatic in every deployment; continuously improving

Two features of the design deserve emphasis for district leaders. First, Vectored Instruction is not a separately purchased module or a setting that must be enabled. It is intrinsic to the Acellus courseware and comes automatically with every deployment of Acellus, in every course, for every student. Second, it improves with use. Because Deficiency Diagnostics is analyzing every response entered by every student, the Acellus courseware development team is alerted to common misconceptions and prerequisite failures across the entire user base and continuously refines items, help videos, and prerequisite bridges. The courses themselves are thus "in a constant state of improvement due to the data collected from Deficiency Diagnostics and Vectored Instruction" (International Academy of Science, 2025, p. 67).

3.7 Where Vectored Instruction sits in the Acellus system

Vectored Instruction is one of several integrated capabilities of the Acellus Gold platform. It works alongside Prism Diagnostics (described above), Tuned Learning and Honors modes that let teachers adjust pacing and challenge for individual students, the Gold Credits and Gold Store incentive system that rewards demonstrated mastery, and the Success Zone reporting that flags satisfactory progress. In the Gold Spark Learning Model, each class period combines roughly forty minutes of personalized digital instruction, during which Vectored Instruction and Prism Diagnostics are finding and filling gaps for every student simultaneously, with twenty minutes of live, teacher-led Special Lessons devoted to discussion, projects, and higher-order thinking (Billings, 2025, 2026). The division of labor is deliberate: the technology handles individualized diagnosis and prerequisite repair, which no single teacher can perform for thirty students at once, and the teacher is freed to do what only a human teacher can do.

4. Outcomes Reported by Districts Using Vectored Instruction

The Acellus Learning System is used in thousands of schools by millions of students across all fifty states, and Vectored Instruction has been operating at scale for years in hundreds of districts. The outcomes summarized here are drawn from the published accounts of district leaders in the *Acellus Journal* and related Acellus publications (Acellus Journal, 2026a, 2026b). They are reported by the districts themselves and are presented as such; Section 9 discusses the value of independent, controlled evaluation going forward.

4.1 Le Grand Union Elementary School District, Merced County, California

Le Grand Union serves a K–8 population that is 96 percent Hispanic/Latino, 80 percent high-poverty, and 40 percent English language learners. Emerging from the pandemic, teachers at every grade were re-teaching content from two and three grades earlier. Superintendent Scott Borba ran a deliberate comparison, placing Acellus Gold with Vectored Instruction in one classroom per grade from second through eighth while parallel classrooms used the previous platform or nothing, holding constant sufficient instructional minutes and active teacher use of the progress dashboard. Within roughly six to eight weeks "the discrepancy between the two classrooms was so evident" that staff asked to expand school-wide. After a full year of school-wide use, Le Grand Union became one of only two schools in all of Merced County to achieve positive academic growth in both English language arts and mathematics relative to *pre-COVID* levels; its students are above the state average in language arts and approaching it in mathematics. Students who entered fourth grade reading at a first-grade level left fifth grade reading on grade level; three-year learning deficits were closed for individual students; and students who might have been referred to special education remained in mainstream classrooms. Teachers now plan and teach grade-level content exclusively: "If you teach fifth grade, you're teaching fifth grade. You're not going back and filling in holes. We have something doing that for us" (Acellus Journal, 2026a, pp. 10–11).

4.2 Community School of Holy Angels (special education)

Holy Angels serves students with severe learning challenges and was among the first schools in the nation to pilot Acellus Gold with Vectored Instruction. In a single school year, mathematics scores rose by up to 40 percent and English scores improved by 24 percent. One student from a non-English-speaking

household, on the autism spectrum and struggling to communicate, began speaking fluently in English after working with the system (Billings, 2026; Acellus, 2026a).

4.3 Lansing Unified School District 469, Kansas (Kansas Heartland Virtual Academy)

Lansing USD 469 built its Kansas Heartland Virtual Academy on the Acellus National CyberSchool model, in which students learn at home on the Acellus platform while remaining fully enrolled in the district. Superintendent Marty Kobza describes Vectored Instruction as "the engine that drives recovery" for students who are behind: "It assesses where a student is on the fly, then goes back to reteach and fill the gaps, all while the student keeps moving forward with grade-level content." The district reports an 18 percent average increase in lesson scores among students being vectored on their retakes; retention in the virtual academy near 100 percent; a rise in the high school graduation rate from about 87 percent to about 96 percent, with Acellus described as "the backbone" of that effort; Writing Tutor gains from roughly 25 percent on first drafts to 95–100 percent on final drafts; and a net financial benefit of approximately \$4,000 per virtual student after program expenses (Acellus Journal, 2026b, pp. 12–13, 23).

4.4 Enrollment and program outcomes

Acton-Agua Dulce Unified School District in California welcomed 180 new students in the first year of its CyberSchool program, an 18 percent enrollment increase (Acellus Journal, 2026a, p. 15). Districts using the system routinely report the pattern noted in the Acellus classroom guide: measurable score increases following Vectored Instruction interventions, visible directly on the teacher dashboard, together with substantial recovery of teacher time otherwise spent hunting for the "one lost concept" for a few students (International Academy of Science, 2025).

District / School	Context	Reported outcomes with Acellus Gold and Vectored Instruction
Le Grand Union ESD, CA	K–8; 96% Hispanic/Latino; 80% high-poverty; 40% ELL; post-COVID regression across all grades	One of two schools in Merced County with positive growth vs. pre-COVID in both ELA and math; above state average in ELA; 3-year individual deficits closed; fewer SPED referrals; teachers teach grade level only
Community School of Holy Angels	Special education; severe learning challenges	Math scores up to +40%; English +24% in one school year; individual communication breakthroughs
Lansing USD 469, KS (Kansas Heartland Virtual Academy)	District-sponsored home-study program on the National CyberSchool model	+18% average lesson-score gain for vectored students; retention ≈100%; graduation rate 87% → 96%; Writing Tutor 25% → 95–100%; ≈\$4,000 net per student
Acton-Agua Dulce USD, CA	CyberSchool program launch	180 new students in year one (+18% enrollment)

Table 1. Summary of district-reported outcomes (Acellus Journal, 2026a, 2026b; Billings, 2026).

4.5 What these outcomes have in common

Three features recur across these reports and are consistent with the mechanism described in Section 3. First, the largest gains appear among students who were furthest below grade level: special-education students, English learners, students in high-poverty schools, and students returning to school after long absences. This is exactly what a prerequisite-repair mechanism predicts, since those students carry the largest and oldest gaps. Second, gains coexist with grade-level instruction rather than replacing it; teachers report *more* time on grade-level content, not less. Third, the effect is visible quickly, in weeks rather than semesters, which is again consistent with just-in-time repair of a specific missing concept as opposed to broad review.

5. The Science of Why Vectored Instruction Works

Vectored Instruction was developed by engineering practice, by getting inside the problem, instrumenting the learning process, and iterating on what the data revealed. It was not derived from a single theory. Yet when the method is set against the research literature of the past six decades, it becomes clear that it operationalizes, in one integrated and automated loop, a set of principles that each carry strong independent evidence. This section reviews those principles. For each, we state the finding, the evidence, and how Vectored Instruction embodies it.

5.1 Learning hierarchies and the primacy of prior knowledge

The principle. Complex skills are built from ordered prerequisite skills; the single strongest determinant of new learning is what the learner already knows.

The evidence. Gagné's (1968) learning-hierarchy studies established that a target skill cannot be reliably acquired until its subordinate skills are in place, and that failure at a higher node is diagnostic of a missing lower node. Ausubel (1968) made prior knowledge the centerpiece of his theory of meaningful learning. Hattie's (2009) synthesis of more than 800 meta-analyses ranks prior achievement ($d \approx 0.55$) and prior ability among the most powerful influences on learning. Siegler et al. (2012) supplied the sharpest longitudinal demonstration: specific elementary prerequisites (fractions and division) predict high-school algebra achievement years later, independent of general ability. Knowledge Space Theory (Doignon & Falmagne, 1985, 1999) formalized the idea mathematically, showing that a domain can be represented as a set of prerequisite-closed *knowledge states*, that a student's state can be identified from a modest number of well-chosen responses, and that the "outer fringe" of that state defines exactly what the student is ready to learn next. This theory underlies the ALEKS system and has been shown to be formally equivalent to modern cognitive diagnostic models (Falmagne, Cosyn, Doignon, & Thiéry, 2006; Heller, Stefanutti, Anselmi, & Robusto, 2015).

How Vectored Instruction embodies it. The entire premise of Vectored Instruction is that failure on a grade-level lesson is most often a symptom of a missing prerequisite. Deficiency Diagnostics is, in Gagné's terms, a procedure for locating the missing node in the hierarchy; Customized Personal Instruction teaches that node; and the return to the original lesson tests whether the hierarchy now holds. Because Acellus courses span kindergarten through twelfth grade in a single integrated system, the prerequisite

structure can be traversed across grade boundaries, which is precisely the traversal that a single-grade textbook, a single-grade teacher, or a single-grade assessment cannot make.

5.2 Mastery learning and the 2 sigma challenge

The principle. Almost all students can master academic material if instruction is organized so that each unit is mastered before the next begins, with corrective instruction supplied when mastery is not demonstrated.

The evidence. Bloom (1968) argued that variation in achievement reflects variation in time and correctives, not fixed aptitude, and that mastery organization could bring 90 percent of students to levels previously reached by 20 percent. Kulik, Kulik, and Bangert-Drowns (1990) meta-analyzed 108 studies and found an average effect of roughly half a standard deviation on examinations, with larger effects for lower-aptitude students and positive effects on retention and attitudes. Bloom (1984) then reported that one-to-one tutoring under mastery conditions produced gains of about two standard deviations, and mastery learning alone about one, framing the "2 sigma problem" of achieving tutoring-level results at scale. Guskey (2007) emphasized a point often lost in practice: the correctives in mastery learning must present the material *differently*, not merely repeat it, and must be followed by a second assessment of mastery.

How Vectored Instruction embodies it. Acellus courses are mastery-based by design: students do not advance until they demonstrate understanding of each concept. Vectored Instruction is the *corrective* stage of the mastery cycle, but with a critical refinement over Bloom's original formulation. Where classic mastery learning re-teaches the same unit with an alternative approach, Vectored Instruction first asks whether the unit itself is the problem, and when it is not, it goes to the prerequisite. Guskey's requirement that the corrective differ from the original instruction is met in the most fundamental way possible: the corrective is a different lesson, from a different grade, addressing the actual root cause. And because the loop is automated and individualized, it approaches the one-to-one condition Bloom identified as the gold standard, but at the scale of an entire student body.

5.3 Cognitive load, schemas, and the zone of proximal development

The principle. Working memory is severely limited; learning new material depends on having well-formed schemas for its components in long-term memory. Instruction that assumes schemas the learner does not have overloads working memory and fails, while instruction that duplicates schemas the learner already has is redundant and can be counterproductive. Learning is most efficient in the zone just beyond what the learner can do alone, with support that is withdrawn as competence grows.

The evidence. Sweller's (1988) cognitive load theory explains why a student without automated fraction schemas cannot follow an algebra lesson: the lesson consumes working-memory capacity on component operations that a prepared student handles automatically. Kalyuga, Ayres, Chandler, and Sweller (2003) documented the *expertise reversal effect*, in which support that helps novices actively hinders more knowledgeable learners, a strong argument for targeting instruction to each student's actual state rather than delivering uniform review to a whole class. Kirschner, Sweller, and Clark (2006) showed that novices need explicit, guided instruction rather than unguided struggle. Vygotsky's (1978) zone of proximal

development and the concept of scaffolding introduced by Wood, Bruner, and Ross (1976) describe the same target from the developmental side; Hattie (2009) reports an average effect for scaffolding of about 0.82.

How Vectored Instruction embodies it. By diagnosing the missing schema and teaching it explicitly before returning the student to the lesson, Vectored Instruction reduces the intrinsic load of the grade-level task to a level the student can manage. Because the intervention is triggered by *that* student's failure and delivered only to that student, it avoids the expertise-reversal penalty that whole-class review imposes on students who already possess the prerequisite. And because the student is returned to the grade-level lesson as soon as the prerequisite is mastered, the scaffold is withdrawn precisely when it is no longer needed. This is why students at Le Grand Union who did not need intervention were not slowed down while those who did were caught up.

5.4 Diagnostic assessment through the analysis of wrong answers

The principle. Wrong answers are not noise. When assessment items are constructed so that specific distractors correspond to specific misconceptions or missing prerequisites, the pattern of a student's errors identifies what the student actually believes and lacks.

The evidence. Sadler (1998) demonstrated that "distractor-driven" multiple-choice tests, whose wrong options are drawn from documented student misconceptions, can diagnose those misconceptions with precision, and Sadler and Sonnert (2016) showed that teachers who know their students' likely misconceptions produce larger learning gains. Treagust (1988) developed the two-tier diagnostic item, and Hestenes, Wells, and Swackhamer (1992) built the Force Concept Inventory on the same logic. Gierl, Bulut, Guo, and Zhang (2017) reviewed the field and confirmed that well-designed distractors carry substantial diagnostic information about student thinking. Diagnostic classification models (Rupp, Templin, & Henson, 2010; Leighton & Gierl, 2007) provide the psychometric machinery to map response patterns to profiles of mastered and unmastered attributes rather than to a single score. Importantly, Gurung et al. (2023) found that teachers predict common wrong answers only inconsistently, which means the diagnostic map must be built from empirical data across many students, not from expert intuition alone.

How Vectored Instruction embodies it. Deficiency Diagnostics is a distractor-driven diagnostic system operating at population scale. Acellus's discovery that large proportions of students choose the *same* wrong answers, and that those choices reveal specific missing concepts, is exactly the phenomenon Sadler documented; the difference is that Acellus observes it continuously across millions of learners and can therefore build and refine the diagnostic map empirically. Prism Diagnostics extends the idea by separating students who share a wrong answer into subgroups with different underlying deficits. The help videos triggered by specific wrong answers, which redirect thinking without giving away the answer, are the instructional counterpart of a two-tier item: they probe the reason behind the response.

5.5 Formative assessment and the power of immediate, task-focused feedback

The principle. Assessment used to adjust instruction while learning is under way produces large gains, particularly for low achievers; feedback is most effective when it is immediate, specific, and focused on the task rather than on the learner.

The evidence. Black and Wiliam's (1998) landmark review reported effect sizes of roughly 0.4 to 0.7 for formative assessment, among the largest for any educational intervention, with the greatest benefits for low-achieving students. Hattie (2009) places feedback near the top of his rankings ($d \approx 0.70$). Kluger and DeNisi (1996), across 607 effect sizes, found that feedback helps on average ($d \approx 0.41$) but that more than a third of feedback interventions actually reduce performance, typically when feedback draws attention to the self ("you are behind") rather than to the task. Shute (2008) synthesized the design guidance: elaborated, task-level, immediate feedback is favored for difficult material and for lower-achieving learners.

How Vectored Instruction embodies it. Every Acellus assessment is formative in the strict sense: its results change what the student is taught next. The feedback loop closes in seconds rather than the days or weeks that elapse when a teacher must grade, analyze, plan, and re-teach. The intervention is entirely task-focused, addressing the concept rather than the student, and the student is not labeled, pulled out, or told they are behind. Kluger and DeNisi's warning about self-directed feedback is thus designed out of the system, which may help explain the motivational effects district leaders describe.

5.6 Retrieval practice and spacing

The principle. Being tested on material strengthens memory more than re-studying it, and reviews spaced over time produce more durable learning than massed review.

The evidence. Roediger and Karpicke (2006) showed that students who were tested after study retained substantially more a week later than those who re-studied, despite lower immediate performance. Cepeda, Vul, Rohrer, Wixted, and Pashler (2008) established that the optimal gap between reviews scales with the desired retention interval.

How Vectored Instruction embodies it. The Acellus mastery cycle is built on frequent retrieval: each concept is assessed immediately after instruction and the student cannot advance without demonstrating it. Vectored Instruction adds a naturally spaced second retrieval of the prerequisite, since a concept first taught in an earlier grade is retrieved and re-taught when it is needed years later, and then applied immediately in the grade-level lesson, which is the strongest possible form of elaborative practice.

5.7 What the intelligent tutoring literature tells us

The principle. Computer-based tutors that track individual knowledge and respond step-by-step can approach the effectiveness of human tutors.

The evidence. VanLehn (2011) reviewed the literature and concluded that step-based intelligent tutoring systems ($d \approx 0.76$) were nearly as effective as human tutoring ($d \approx 0.79$) and far more effective than answer-based computer instruction ($d \approx 0.31$). Ma, Adesope, Nesbit, and Liu (2014) meta-analyzed 107 effect sizes and found that intelligent tutoring outperformed teacher-led large-group instruction ($g \approx 0.42$)

and non-adaptive computer instruction ($g \approx 0.57$) and was statistically indistinguishable from human tutoring. Kulik and Fletcher (2016) reported a median effect of 0.66 standard deviations across 50 controlled evaluations. The literature also contains a caution: Steenbergen-Hu and Cooper (2013) found small effects for some K–12 mathematics tutors, especially for low achievers, which underscores that the *design* of the adaptivity matters. Systems that merely adjust difficulty within the current topic are not the same as systems that diagnose and repair prerequisites. Escueta, Nickow, Oreopoulos, and Quan (2020), reviewing randomized trials of educational technology, likewise found that adaptive computer-assisted learning, particularly in mathematics, shows the most consistent positive effects of any category of educational technology.

How Vectored Instruction embodies it. Acellus with Vectored Instruction is a step-based system in VanLehn's sense (it responds at the level of individual concepts and items, not merely at the level of the final answer), and it adds a dimension most tutoring systems lack: cross-grade traversal of the prerequisite hierarchy. The evidence that low achievers gain least from tutors that stay within the current topic, and most from mastery learning and formative feedback that reach the root cause, is precisely the pattern that motivated the development of Vectored Instruction and that the district outcomes in Section 4 reflect.

5.8 Acceleration, not remediation; grade level, not retention

The principle. Students learn more when they are kept on grade-level content and given just-in-time support on the specific prerequisites they lack, than when they are held back to below-grade work or retained in grade.

The evidence. TNTP's (2018) study of nearly 4,000 students found that students spent more than 500 hours a year on assignments below grade level, met grade-level standards only 17 percent of the time, and that classrooms with more grade-appropriate work gained the equivalent of about two additional months of learning. TNTP and Zearn (2021), analyzing more than two million students, found that students who received grade-level content with just-in-time scaffolding on prerequisites completed 27 percent more grade-level lessons (49 percent more in schools serving predominantly students of color) and struggled *less* than matched peers who were remediated with below-grade content first. Randomized evidence from higher education points the same way: students placed directly into college-level statistics with embedded ("corequisite") support passed at 48 percent versus 31–36 percent for those routed through prerequisite remedial algebra, and they earned more credits and persisted longer (Logue, Watanabe-Rose, & Douglas, 2016). Meanwhile, Jimerson's (2001) meta-analysis of grade retention found no academic or socio-emotional benefit relative to promotion.

How Vectored Instruction embodies it. Vectored Instruction is the purest available implementation of "accelerate, don't remediate." The student never leaves the grade-level course. The prerequisite is supplied *just in time*, at the exact moment it is needed to unlock the current lesson, and only the specific prerequisite that is missing, not a whole year of below-grade review "just in case." The design therefore honors both the district's obligation to keep students at grade level and the research consensus that this is also what is best for the student.

5.9 Solving the scale problem

The principle. The benefits of individualized diagnosis and tutoring are well established; the barrier has always been delivering them to every student.

The evidence. Nickow, Oreopoulos, and Quan (2024) confirm the effectiveness of tutoring; Kraft, Schueler, and Falken (2024) show that effects fall to roughly 0.16 standard deviations once programs exceed a thousand students, because skilled tutors are scarce, expensive, and hard to schedule. Teacher shortages exceeding 425,000 positions (Learning Policy Institute, 2026) make the constraint sharper still.

How Vectored Instruction embodies it. Because diagnosis and prerequisite repair are automated, Vectored Instruction delivers to every student, in every course, at every moment of need, the two components of tutoring that scale worst: figuring out precisely what the student is missing, and teaching exactly that. It thereby frees human teachers to concentrate on mentoring, discussion, projects, and encouragement, which is the division of labor that the Gold Spark model formalizes and that superintendents describe as giving teachers back to students.

5.10 Motivation, self-efficacy, and the message the student hears

Finally, the affective dimension deserves mention. D'Mello and Graesser (2012) mapped the sequence by which unresolved confusion decays into frustration and then boredom during complex learning; the moment of being stuck is the moment when engagement is won or lost. Traditional responses to failure (re-teaching louder, pull-out remediation, labels such as "behind" or "unmotivated") arrive late and direct attention to the self, precisely the conditions under which Kluger and DeNisi (1996) found feedback to be harmful. Vectored Instruction intervenes at the point of confusion, before frustration sets in, and it intervenes in a way that is invisible and non-stigmatizing. Combined with the Gold Credits earned for demonstrated mastery, the student experiences a run of successes where previously there was a wall. As Superintendent Borba observed, he does not hear students or families complaining about Acellus as he has with other platforms; and at Holy Angels, students "the system said couldn't keep up" began keeping up (Acellus Journal, 2026a; Billings, 2026).

Research principle	Key evidence	Where it appears in Vectored Instruction
Learning hierarchies / prior knowledge	Gagné 1968; Ausubel 1968; Siegler et al. 2012; Doignon & Falmagne 1985	Failure treated as symptom of missing prerequisite; cross-grade traversal to the missing node
Mastery learning	Bloom 1968, 1984; Kulik et al. 1990; Guskey 2007	No advance without mastery; corrective is a different lesson addressing the root cause; re-assessment after repair
Cognitive load; expertise reversal; ZPD	Sweller 1988; Kalyuga et al. 2003; Kirschner et al. 2006; Wood et al. 1976	Explicit teaching of the missing schema only to the student who lacks it; scaffold withdrawn on return to lesson

Research principle	Key evidence	Where it appears in Vectored Instruction
Distractor-driven diagnosis	Sadler 1998; Treagust 1988; Gierl et al. 2017; Gurung et al. 2023	Deficiency Diagnostics and Prism Diagnostics; empirically built wrong-answer map; targeted help videos
Formative assessment; feedback	Black & Wiliam 1998; Kluger & DeNisi 1996; Shute 2008	Instant, task-focused, non-labeling response; loop closes in seconds
Retrieval practice; spacing	Roediger & Karpicke 2006; Cepeda et al. 2008	Mastery checks after every concept; spaced re-retrieval of prerequisites
Intelligent tutoring effectiveness	VanLehn 2011; Ma et al. 2014; Kulik & Fletcher 2016	Step-based adaptivity plus prerequisite repair; strongest for lowest achievers
Accelerate, don't remediate; avoid retention	TNTP 2018; TNTP & Zearn 2021; Logue et al. 2016; Jimerson 2001	Student never leaves grade-level course; just-in-time, not just-in-case
Tutoring at scale	Bloom 1984; Nickow et al. 2024; Kraft et al. 2024	Automated diagnosis and repair for every student; teachers freed for human work

Table 2. Research foundations of Vectored Instruction.

6. Why Vectored Instruction Is Especially Crucial for Students Studying at Home

Homeschooling and district-sponsored independent study have become the fastest-growing segment of American education. Roughly three million students, approaching six percent of the K–12 population, are now homeschooled, about double the pre-pandemic share; growth averaged nearly five percent across reporting states in 2024–25, with ten states at record highs (Johns Hopkins Homeschool Hub, 2025; EdChoice, 2026). Over the same period public school enrollment fell by 1.4 million students, and because state funding follows per-pupil counts, the fiscal consequences for districts are substantial (FutureEd, 2026). Districts that offer supported home-study programs are responding to a real demand.

But the student studying at home faces a specific and serious risk that the classroom student does not. In a classroom, when a student becomes stuck, a professional teacher is physically present. The teacher may not be able to diagnose the missing prerequisite precisely, but the teacher can notice the struggle, ask a question, try another explanation, or arrange help. At home, no such professional is present. The parent may be highly motivated but is seldom trained to recognize that a tenth grader's algebra failure is really a fifth-grade fraction failure, and even less able to teach the prerequisite. Online programs that simply deliver content over video, in Superintendent Kobza's phrase, "the same methodology over video or Zoom," leave the student alone with the wall. The predictable result is the one districts have long observed: students who leave to study at home fall behind, and when they return years later the district is expected to catch them up (International Academy of Science, 2026).

Vectored Instruction changes this equation fundamentally. It supplies, automatically and at the moment of need, exactly the function that the absent teacher would otherwise have to provide: recognizing that

the student is stuck, determining why, teaching the missing piece, and returning the student to the lesson. Because it is embedded in the courseware, it is present at the kitchen table at nine o'clock at night as reliably as it is in a classroom at nine o'clock in the morning. For an independent-study student, then, Vectored Instruction is not an enhancement; it is the safety net that makes home study academically defensible. It is why the Kansas Heartland Virtual Academy can report near-total retention alongside 18 percent lesson-score gains for vectored students, and why students who return from that program to the traditional building "arrive confident and on track," with "not a single problem" in the transition (Acellus Journal, 2026b, p. 13). The parents' role in such a program is to encourage and stay involved, not to be the full-time instructor; Vectored Instruction, the supervising district teacher, and the Acellus support team carry the instructional load.

For district leaders considering any independent-study or virtual option, this is the decisive question to ask of a platform: *when the student gets stuck at home, what happens?* If the answer is "the parent helps" or "the student emails a teacher and waits," the program will produce the outcomes that have historically given virtual learning a poor reputation. If the answer is "the system diagnoses the missing foundation and repairs it, immediately, and the teacher sees it happen on the dashboard," the program has a mechanism to keep students on grade level while they learn away from the building.

7. Using Vectored Instruction in Your Teaching

Vectored Instruction operates automatically, but the districts reporting the strongest results are those in which teachers actively work *with* it. This section distills the practices those districts describe (Acellus Journal, 2026a, 2026b; International Academy of Science, 2025) together with the research reviewed above into practical guidance for educators. The first subsection addresses teachers using Acellus; the second addresses the many educators who do not yet have the technology but wish to apply its principles.

7.1 For teachers and administrators using Acellus

1. **Protect the instructional minutes.** Le Grand Union held two variables constant in its comparison classrooms: sufficient time on the platform, and teachers actively engaging with the dashboard. Vectored Instruction cannot diagnose or repair a gap in a student who is not working through the lessons. Establish a daily block (the forty-minute digital phase of the Gold Spark model is the reference design) and protect it.
2. **Watch the Vectored Instruction panel daily.** The Active Teacher Interface shows every intervention, the concept diagnosed, the bridging lesson used, and before-and-after scores. Use it as a formative-assessment instrument. A cluster of interventions on the same prerequisite is a signal that a short whole-class Special Lesson on that concept would pay off; a student with repeated interventions across concepts may need Tuned Learning mode or additional human attention.
3. **Teach grade level; let the system fill the gaps.** The most transformative change reported at Le Grand Union was that teachers stopped re-teaching earlier grades and planned exclusively for

their own. Trust the mechanism, verify it on the dashboard, and redirect the recovered time to discussion, projects, and the students the dashboard flags.

4. **Use Special Lessons for what only humans can do.** Because content delivery and prerequisite repair are handled individually, the teacher-led portion of the period is free for the higher-order work the research associates with deep learning: discussion, argument, project-based learning, hands-on demonstration, and personal encouragement. Use the dashboard data to make those lessons targeted rather than generic.
5. **Do not label the vectored student.** The design of Vectored Instruction deliberately avoids telling the student they are "behind." Preserve that. Talk about concepts, not deficits; celebrate the "after" score; and let Gold Credits do their motivational work.
6. **Enable Tuned Learning and Honors deliberately.** For students who are consistently outside the Success Zone despite interventions, Tuned Learning slows the pace and increases the step-by-step support, including Vectored Instruction; for students who are far ahead, Honors mode adds challenge. Both keep the whole class moving toward the goal of 100 percent of students in the Success Zone.
7. **Give it time, and expect the first evidence in weeks.** Le Grand Union saw a clear discrepancy between comparison classrooms in six to eight weeks and school-wide results within a year. Plan the evaluation window accordingly and share the dashboard evidence with staff and the board.

7.2 Applying the principles without the platform

Teachers who do not have access to Acellus can nevertheless move their practice toward the Vectored Instruction model. The following are the core moves, each grounded in the research of Section 5.

- **Reframe every failure as a diagnostic question.** When a student misses a grade-level item, ask first "which prerequisite would explain this?" rather than "how do I re-teach this lesson?" Keep a prerequisite map for your course (Section 8.1) so that the question can be answered quickly.
- **Write diagnostic distractors.** Following Treagust (1988), Sadler (1998), and Haladyna, Downing, and Rodriguez (2002), build wrong-answer options from the errors students actually make, and record which distractor each student chooses. Over a few years the pattern will tell you which prerequisites your incoming students most often lack.
- **Prepare bridge lessons in advance.** For the five to ten prerequisites that most often derail your course, prepare short, self-contained explanations, worked examples, and practice items (Renkl, 2014), ideally as short videos students can watch independently. Deliver them to individual students at the moment of need, not to the whole class in advance.
- **Keep the student in the grade-level task.** Supply the bridge and return the student to the original problem the same day. Avoid pull-out remediation or extended below-grade review; the evidence favors just-in-time support inside grade-level work (TNTP & Zearn, 2021).

- **Close the loop with a re-assessment.** A bridge lesson is complete only when the student demonstrates the prerequisite and then succeeds on the original item (Guskey, 2007).
- **Make feedback about the task, immediate, and private.** Follow Shute (2008) and Kluger and DeNisi (1996): specific, elaborated, task-level feedback delivered quickly; nothing that draws attention to the student's standing relative to peers.

8. Applying Vectored Instruction Principles in Textbooks and Learning Resources

Most instructional materials, print and digital, are organized as a linear sequence for a single grade. They implicitly assume that every prerequisite from earlier grades is in place, and when that assumption fails, the resource has nothing to offer. Developers of textbooks, courseware, and open educational resources can build the logic of Vectored Instruction into their products. The following design principles apply whether the medium is paper, PDF, or software.

8.1 Build and publish the prerequisite map

The foundation of any vectored resource is an explicit map of which concepts depend on which earlier concepts, across grade boundaries. Curriculum research supplies starting points: the learning progressions literature (Heritage, 2008; Corcoran, Mosher, & Rogat, 2009), Zimba's (2012) dependency "wiring diagram" of the mathematics standards, and the interactive Coherence Map built on it (Student Achievement Partners, 2015), which draws explicit "builds on" arrows from each standard to its prerequisites in earlier grades. Knowledge Space Theory (Doignon & Falmagne, 1999) offers a rigorous formalism for the map, and data-driven methods can discover or validate prerequisite relations from student performance (Scheines, Silver, & Goldin, 2014; Chen, Lu, Zheng, & Pian, 2018). Whatever the source, the map must span grades: a seventh-grade ratio unit must know that it depends on fifth-grade fractions and third-grade multiplication.

8.2 Design assessments that diagnose, not merely score

Every end-of-lesson check should be written so that specific wrong answers point to specific prerequisites in the map. Use two-tier items where practical (answer plus reason; Treagust, 1988), draw distractors from documented student errors rather than expert guesses (Gurung et al., 2023), and pilot items to confirm that each distractor is chosen for the reason intended (Gierl et al., 2017). In print, this can be as simple as a coded answer key: "If the student chose (b), see Bridge 4.2." In software, it is the routing table for automatic vectoring.

8.3 Author bridge lessons, not review chapters

For each prerequisite that the map identifies as a frequent failure point, provide a compact, self-contained *bridge lesson*: a brief explanation, worked examples with fading (Renkl, 2014), a few practice items with feedback, and a mastery check, followed by an explicit return pointer to the grade-level lesson. Bridge lessons should be written for the older student who is encountering the concept a second time, not simply

copied from the earlier grade's text. They belong at the point of need (a margin pointer, a QR code, a hyperlink, an automatic branch), not in a front-of-book review chapter that assumes the reader knows what they do not know.

8.4 Keep the grade-level spine intact

A vectored resource never removes the student from the grade-level sequence. Bridges are excursions with a guaranteed return path. Backward design (Wiggins & McTighe, 2005) is the appropriate authoring discipline: define the grade-level outcome and its evidence first, then determine which prerequisites the evidence depends on, then write the bridges. Multi-tiered support frameworks (Fuchs & Fuchs, 2006) and the What Works Clearinghouse guidance on mathematics intervention (Fuchs et al., 2021) supply criteria for the intensity and design of the bridges themselves, and Universal Design for Learning (CAST, 2024) should govern their accessibility.

8.5 Instrument the resource and improve it

Digital resources should log which items are missed, which distractors are chosen, which bridges are triggered, and whether the return to the grade-level item succeeds. That data validates the prerequisite map, reveals bridges that do not work, and identifies missing bridges. This is the continuous-improvement loop that Deficiency Diagnostics provides at Acellus and that learning-engineering practice recommends generally (Koedinger et al., 2010; Prihar et al., 2022). Open educational resources are particularly well suited to this model, since bridge lessons can be remixed and shared across products at low cost without loss of learning outcomes (Hilton, 2020).

Design checklist for a vectored learning resource

1. Cross-grade prerequisite map, published with the resource.
2. Diagnostic items whose distractors route to specific prerequisites.
3. Compact bridge lessons at the point of need, each with a mastery check and a return pointer.
4. Grade-level spine preserved; no removal to below-grade sequences.
5. Logging and review of item, distractor, and bridge performance to refine the map.

9. Enhancing Vectored Instruction Going Forward

Vectored Instruction is a mature technology in daily use, but the research frontier suggests several directions in which the next generation can become faster, more precise, and more broadly applicable. We group them under diagnosis, instruction, retention, teacher partnership, and evidence.

9.1 Sharper diagnosis

- **Predictive vectoring.** Bayesian Knowledge Tracing (Corbett & Anderson, 1995), Deep Knowledge Tracing (Piech et al., 2015), and transformer-based models (Choi et al., 2020; Abdelrahman,

Wang, & Nunes, 2023) estimate the probability that a student has mastered each skill from their full interaction history. Integrated with the prerequisite map, such models could predict a coming failure before it occurs and vector proactively, delivering the bridge just ahead of the lesson that would otherwise stall. Prerequisite-driven variants (Chen et al., 2018) are directly suited to this use.

- **Data-discovered prerequisite structure.** Causal-discovery and prerequisite-learning methods (Scheines et al., 2014; Pan, Li, Li, & Tang, 2017) can mine the millions of Acellus response records to confirm known dependencies, discover unexpected ones, and identify bridges that should exist but do not.
- **Distinguishing gaps from disengagement.** Detectors of wheel-spinning (Beck & Gong, 2013), of gaming the system (Baker, Corbett, Roll, & Koedinger, 2008), and of affective states such as confusion and frustration (D'Mello & Graesser, 2012) can help the system decide whether a string of errors reflects a missing prerequisite (vector), a motivational lapse (encourage, alert the teacher), or guessing (slow down and re-engage), each of which calls for a different response.

9.2 Richer instruction

- **Generative bridges with guardrails.** Large language models can generate individualized explanations, additional worked examples, and Socratic hints on demand. The evidence is now clear about how they must be deployed: an unguarded chatbot raised practice scores but *lowered* unassisted exam performance by 17 percent, while a version with pedagogical guardrails removed the harm (Bastani et al., 2025); a carefully designed AI tutor in a university physics course produced more than twice the learning of active-learning classes in less time (Kestin, Miller, Klales, Milbourne, & Ponti, 2025); and conditioning LLM remediation on expert decisions about error type and strategy improves its quality (Wang, Zhang, Robinson, Loeb, & Demszky, 2024a). The natural role for generative AI in Vectored Instruction is therefore to *augment* the diagnosed, curated bridge lesson with adaptive dialogue and extra examples, never to replace diagnosis or to give answers away, in keeping with the "humans in the loop" guidance of the U.S. Department of Education (2023) and UNESCO (Miao & Holmes, 2023).
- **Multimodal input.** Handwritten and spoken responses would let the diagnostic engine see a student's work, not just the final choice, sharpening diagnosis in mathematics especially; the field is early but advancing (Rodrigues et al., 2024).
- **Bridges for special populations.** Personalization for English learners and for students with dyslexia, dyscalculia, and other learning differences is a growing research area (Salas-Pilco et al., 2024). Bridge lessons can be authored in multiple modalities and languages, and Deficiency Diagnostics can learn which variants work best for which learners.

9.3 Durable retention

- **Scheduled re-retrieval of repaired prerequisites.** Once a prerequisite has been rebuilt, spacing research (Cepeda et al., 2008) and trainable spaced-repetition models (Settles & Meeder, 2016)

can schedule brief, well-timed reviews so that the repair holds across the year and into the next grade, and the system can verify that it has.

9.4 A stronger teacher partnership

- **Explainable, actionable dashboards.** Real-time teacher-awareness tools that direct attention to the students who most need it have been shown to improve learning, particularly for lower-prior-knowledge students (Holstein, McLaren, & Aleven, 2018). The Vectorized Instruction panel can evolve toward recommending, from the day's intervention pattern, the single most valuable Special Lesson a teacher could give.
- **AI support for the supervising teacher.** In tutoring settings, an LLM "copilot" that suggests expert moves to human tutors raised student mastery, most for the least experienced tutors, at very low cost (Wang, Ribeiro, Robinson, Loeb, & Demszky, 2024b). An analogous copilot for CyberSchool supervising teachers and parent-support specialists could translate dashboard data into concrete conversation prompts.
- **Interoperability and privacy.** Adopting event-data standards such as 1EdTech Caliper or xAPI would allow prerequisite diagnostics to draw on evidence from other tools a district uses, and would let districts carry learner records forward, all within FERPA and COPPA obligations (1EdTech Consortium, 2020; U.S. Department of Education, PTAC).

9.5 Evidence and continuous improvement

- **Independent, controlled evaluation.** The district-reported outcomes in Section 4 are compelling and consistent with the mechanism, and the classroom-level comparison at Le Grand Union is a model of practical local evaluation. The next step is independent quasi-experimental and randomized studies of Vectorized Instruction that meet the evidence tiers of the Every Student Succeeds Act, of the kind that have been conducted for other adaptive platforms (Cook, Grant, & Ross, 2022; Hashim, 2024). Acellus's scale and instrumentation make such studies unusually feasible.
- **Embedded experimentation.** Platform-embedded A/B testing (Prihar et al., 2022) can compare alternative bridge lessons, alternative triggering thresholds, and alternative feedback designs at scale, so that every element of Vectorized Instruction is continuously optimized against measured learning, in the manner of the Knowledge-Learning-Instruction framework (Koedinger, Corbett, & Perfetti, 2012).

None of these directions changes the core insight. The student who is stuck is almost always missing something specific from earlier; find it, fix it, and bring the student forward. Everything above is in service of finding it sooner, fixing it better, and proving it more rigorously.

10. Conclusion

The grade-level dilemma has constrained American education for generations: students must be taught the current grade, but many cannot learn it because of what they missed in earlier grades, and neither

retention nor generic remediation helps. Vectored Instruction resolves the dilemma by doing, automatically and for every student, what the research on prerequisite knowledge, mastery learning, cognitive load, diagnostic assessment, formative feedback, and tutoring has long said should be done and what no classroom has had the capacity to do: detect the specific missing foundation, rebuild it at the moment of need, and return the student to the grade-level lesson without ever removing them from it. Districts using it report the outcomes that the science predicts, with the largest gains among students who were furthest behind. It comes with every deployment of Acellus. It is indispensable for students learning at home, where no professional teacher is present. And it is now supported by a research base rich enough to guide both its wider adoption and its continued improvement.

For thought leaders in the field, the invitation is to build on the principles described here in research and in new instructional resources. For district superintendents, the invitation is more direct. Every day a child remains stuck on a missing foundation, the gap grows. The tools to find the hole and fill it exist, they are working in schools across the country, and they are available to every district. The remaining question is the one that has always mattered most: whether we will use them.

Vectored Instruction® and Prism Diagnostics® are registered trademarks of the International Academy of Science / Acellus. The author is the creator of Acellus and the originator of Vectored Instruction, and serves as President of the International Academy of Science, which publishes the Journal of Science; this interest is disclosed in accordance with the Journal's Code of Ethics.

About the Author

Dr. Roger E. Billings is the creator of Acellus and the originator of Vectored Instruction. He is President of the International Academy of Science, the Kansas City–area nonprofit research and education organization he co-founded, which developed and operates the Acellus Learning System, Acellus Academy, and the International Journal of Science. Dr. Billings introduced Acellus in 2001 to answer a single question that has driven his work ever since: *Why can't this child learn?* Two decades of instrumenting the learning process across millions of students led to the discovery that struggling students overwhelmingly fail because of specific, identifiable gaps in earlier foundational knowledge, and to the invention of Deficiency Diagnostics and Vectored Instruction, which find those gaps and repair them automatically while keeping students on grade level. He is also one of the architects of Prism Diagnostics, the Acellus Gold platform, the Gold Spark Learning Model, and the National CyberSchool program.

Trained as an engineer and scientist, Dr. Billings is widely known as a pioneer of hydrogen energy. As a high school senior in 1966 he converted a Model A Ford to run on hydrogen and won top honors at the International Science Fair; he went on to found Billings Energy Corporation, to demonstrate hydrogen-fueled buses, automobiles, and the first hydrogen-powered home, and to advance metal-hydride hydrogen storage. In computing he holds early patents in distributed data processing and has founded companies in networking, cryptography, and security. He has been recognized as one of the U.S. Jaycees' Ten Outstanding Young Americans and as a Champion of Science in America, and he holds a doctorate in research from the Institute of Science and Technology. A father and grandfather, he describes his motivation simply: he built Acellus because he wanted something better for his own children and then

realized that every child deserves the same. He speaks and writes frequently for educators nationwide, and can be reached through the International Academy of Science at science.edu.

References

- 1EdTech Consortium. (2020). *Caliper Analytics v1.2 specification*. <https://www.imsglobal.org/spec/caliper/v1p2>
- Abdelrahman, G., Wang, Q., & Nunes, B. (2023). Knowledge tracing: A survey. *ACM Computing Surveys*, 55(11), Article 224. <https://doi.org/10.1145/3569576>
- Acellus Journal. (2026a). Volume 1, Issue 1. International Academy of Science. (Includes "Schools that Shine: Le Grand Union Elementary," "Launching the Acellus National Cyberschool Program," and the feature adaptation of Dr. Roger Billings' keynote, "Vectored Instruction: Finding the Hole and Filling It.")
- Acellus Journal. (2026b). Volume 1, Issue 2. International Academy of Science. (Includes "Schools that Shine: Lansing USD 469 / Kansas Heartland Virtual Academy" and the scholarly feature "Acellus National Cyberschool.")
- Ausubel, D. P. (1968). *Educational psychology: A cognitive view*. Holt, Rinehart & Winston.
- Baker, R. S. J. d., Corbett, A. T., Roll, I., & Koedinger, K. R. (2008). Developing a generalizable detector of when students game the system. *User Modeling and User-Adapted Interaction*, 18, 287–314.
- Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakçı, Ö., & Mariman, R. (2025). Generative AI without guardrails can harm learning: Evidence from high school mathematics. *Proceedings of the National Academy of Sciences*, 122(26), e2422633122.
- Beck, J. E., & Gong, Y. (2013). Wheel-spinning: Students who fail to master a skill. In *Artificial Intelligence in Education (AIED 2013)*, LNCS 7926 (pp. 431–440). Springer.
- Billings, R. E. (2025). *Acellus Gold Spark Learning Model: Full paper*. International Academy of Science.
- Billings, R. E. (2026). *How Acellus is waking up education*. Keynote address, Acellus National Administrators Conference; adapted in Acellus Journal, 1(1), 19–23.
- Black, P., & Wiliam, D. (1998). Inside the black box: Raising standards through classroom assessment. *Phi Delta Kappan*, 80(2), 139–148.
- Bloom, B. S. (1968). Learning for mastery. *Evaluation Comment*, 1(2), 1–12.
- Bloom, B. S. (1984). The 2 sigma problem: The search for methods of group instruction as effective as one-to-one tutoring. *Educational Researcher*, 13(6), 4–16.
- CAST. (2024). *Universal Design for Learning guidelines, version 3.0*. <https://udlguidelines.cast.org/>
- Cepeda, N. J., Vul, E., Rohrer, D., Wixted, J. T., & Pashler, H. (2008). Spacing effects in learning: A temporal ridgeline of optimal retention. *Psychological Science*, 19(11), 1095–1102.
- Chen, P., Lu, Y., Zheng, V. W., & Pian, Y. (2018). Prerequisite-driven deep knowledge tracing. In *2018 IEEE International Conference on Data Mining* (pp. 39–48).
- Choi, Y., et al. (2020). Towards an appropriate query, key, and value computation for knowledge tracing. In *Proceedings of the Seventh ACM Conference on Learning @ Scale*.
- Cook, M. A., Grant, A. A., & Ross, S. M. (2022). *The impacts of i-Ready Personalized Instruction on student math achievement in multiple school districts*. Johns Hopkins Center for Research and Reform in Education.
- Corbett, A. T., & Anderson, J. R. (1995). Knowledge tracing: Modeling the acquisition of procedural knowledge. *User Modeling and User-Adapted Interaction*, 4, 253–278.

- Corcoran, T., Mosher, F. A., & Rogat, A. (2009). *Learning progressions in science: An evidence-based approach to reform* (CPRE Research Report RR-63). Consortium for Policy Research in Education.
- D'Mello, S., & Graesser, A. (2012). Dynamics of affective states during complex learning. *Learning and Instruction*, 22(2), 145–157.
- Doignon, J.-P., & Falmagne, J.-C. (1985). Spaces for the assessment of knowledge. *International Journal of Man-Machine Studies*, 23(2), 175–196.
- Doignon, J.-P., & Falmagne, J.-C. (1999). *Knowledge spaces*. Springer.
- EdChoice. (2026, May 27). *How homeschooling became the fastest-growing education model*. <https://www.edchoice.org/>
- Escueta, M., Nickow, A. J., Oreopoulos, P., & Quan, V. (2020). Upgrading education with technology: Insights from experimental research. *Journal of Economic Literature*, 58(4), 897–996.
- Fahle, E. M., Kane, T. J., Reardon, S. F., & Staiger, D. O. (2024). *Education Recovery Scorecard*. Harvard Center for Education Policy Research and Stanford Educational Opportunity Project. <https://cepr.harvard.edu/education-recovery-scorecard>
- Falmagne, J.-C., Cosyn, E., Doignon, J.-P., & Thiéry, N. (2006). The assessment of knowledge, in theory and in practice. In *Formal Concept Analysis*, LNCS 3874 (pp. 61–79). Springer.
- Fuchs, D., & Fuchs, L. S. (2006). Introduction to Response to Intervention: What, why, and how valid is it? *Reading Research Quarterly*, 41(1), 93–99.
- Fuchs, L. S., et al. (2021). *Assisting students struggling with mathematics: Intervention in the elementary grades* (WWC 2021006). What Works Clearinghouse, Institute of Education Sciences.
- FutureEd. (2026). *K-12 public school enrollment declines, explained*. Georgetown University. <https://www.future-ed.org/>
- Gagné, R. M. (1968). Learning hierarchies. *Educational Psychologist*, 6(1), 1–9.
- Gierl, M. J., Bulut, O., Guo, Q., & Zhang, X. (2017). Developing, analyzing, and using distractors for multiple-choice tests in education: A comprehensive review. *Review of Educational Research*, 87(6), 1082–1116.
- Gurung, A., Baral, S., Vanacore, K. P., McReynolds, A. A., Kreisberg, H., Botelho, A. F., Shaw, S. T., & Heffernan, N. T. (2023). Identification, exploration, and remediation: Can teachers predict common wrong answers? In *LAK23: 13th International Learning Analytics and Knowledge Conference*. ACM.
- Guskey, T. R. (2007). Closing achievement gaps: Revisiting Benjamin S. Bloom's "Learning for Mastery." *Journal of Advanced Academics*, 19(1), 8–31.
- Haladyna, T. M., Downing, S. M., & Rodriguez, M. C. (2002). A review of multiple-choice item-writing guidelines for classroom assessment. *Applied Measurement in Education*, 15(3), 309–333.
- Hashim, S. (2024). Measuring the efficacy of Zearn Math in Louisiana. *AERA Open*.
- Hattie, J. (2009). *Visible learning: A synthesis of over 800 meta-analyses relating to achievement*. Routledge.
- Heller, J., Stefanutti, L., Anselmi, P., & Robusto, E. (2015). On the link between cognitive diagnostic models and knowledge space theory. *Psychometrika*, 80(4), 995–1019.
- Heritage, M. (2008). *Learning progressions: Supporting instruction and formative assessment*. Council of Chief State School Officers.
- Hestenes, D., Wells, M., & Swackhamer, G. (1992). Force Concept Inventory. *The Physics Teacher*, 30(3), 141–158.
- Hilton, J. (2020). Open educational resources, student efficacy, and user perceptions: A synthesis of research published between 2015 and 2018. *Educational Technology Research and Development*, 68, 853–876.

- Holstein, K., McLaren, B. M., & Aleven, V. (2018). Student learning benefits of a mixed-reality teacher awareness tool in AI-enhanced classrooms. In *Artificial Intelligence in Education (AIED 2018)*, LNCS 10947 (pp. 154–168). Springer.
- International Academy of Science. (2025). *Creating the 21st-century classroom: A guide to teaching with Acellus Gold*. International Academy of Science.
- International Academy of Science. (2026). *National Cyberschool backgrounder*. International Academy of Science.
- Jimerson, S. R. (2001). Meta-analysis of grade retention research: Implications for practice in the 21st century. *School Psychology Review*, 30(3), 420–437.
- Johns Hopkins Homeschool Hub. (2025). *Homeschool growth 2024–2025*. Johns Hopkins School of Education. <https://education.jhu.edu/edpolicy/policy-research-initiatives/homeschool-hub/>
- Kalyuga, S., Ayres, P., Chandler, P., & Sweller, J. (2003). The expertise reversal effect. *Educational Psychologist*, 38(1), 23–31.
- Kestin, G., Miller, K., Klaes, A., Milbourne, T., & Ponti, G. (2025). AI tutoring outperforms in-class active learning: An RCT introducing a novel research-based design in an authentic educational setting. *Scientific Reports*, 15, 17458.
- Kirschner, P. A., Sweller, J., & Clark, R. E. (2006). Why minimal guidance during instruction does not work. *Educational Psychologist*, 41(2), 75–86.
- Kluger, A. N., & DeNisi, A. (1996). The effects of feedback interventions on performance: A historical review, a meta-analysis, and a preliminary feedback intervention theory. *Psychological Bulletin*, 119(2), 254–284.
- Koedinger, K. R., Baker, R. S. J. d., Cunningham, K., Skogsholm, A., Leber, B., & Stamper, J. (2010). A data repository for the EDM community: The PSLC DataShop. In *Handbook of Educational Data Mining*. CRC Press.
- Koedinger, K. R., Corbett, A. T., & Perfetti, C. (2012). The Knowledge-Learning-Instruction framework: Bridging the science-practice chasm to enhance robust student learning. *Cognitive Science*, 36(5), 757–798.
- Kraft, M. A., Schueler, B. E., & Falken, G. (2024). *What impacts should we expect from tutoring at scale? Exploring meta-analytic generalizability* (EdWorkingPaper). Annenberg Institute, Brown University.
- Kulik, C.-L. C., Kulik, J. A., & Bangert-Drowns, R. L. (1990). Effectiveness of mastery learning programs: A meta-analysis. *Review of Educational Research*, 60(2), 265–299.
- Kulik, J. A., & Fletcher, J. D. (2016). Effectiveness of intelligent tutoring systems: A meta-analytic review. *Review of Educational Research*, 86(1), 42–78.
- Learning Policy Institute. (2026, July). *State teacher shortages 2026: Fact sheet*. <https://learningpolicyinstitute.org/>
- Leighton, J. P., & Gierl, M. J. (Eds.). (2007). *Cognitive diagnostic assessment for education: Theory and applications*. Cambridge University Press.
- Lewis, K., & Kuhfeld, M. (2024). *Recovery still elusive: 2023–24 student achievement highlights persistent achievement gaps and a long road ahead*. NWEA.
- Logue, A. W., Watanabe-Rose, M., & Douglas, D. (2016). Should students assessed as needing remedial mathematics take college-level quantitative courses instead? A randomized controlled trial. *Educational Evaluation and Policy Analysis*, 38(3), 578–598.
- Ma, W., Adesope, O. O., Nesbit, J. C., & Liu, Q. (2014). Intelligent tutoring systems and learning outcomes: A meta-analysis. *Journal of Educational Psychology*, 106(4), 901–918.
- Miao, F., & Holmes, W. (2023). *Guidance for generative AI in education and research*. UNESCO.

- National Assessment Governing Board. (2025a, January). *The Nation's Report Card 2024: 10 takeaways from the 2024 NAEP results*. <https://www.nagb.gov/>
- National Assessment Governing Board. (2025b, September). *Top 5 takeaways from 12th-grade NAEP math and reading results*. <https://www.nagb.gov/>
- Nickow, A., Oreopoulos, P., & Quan, V. (2024). The promise of tutoring for PreK–12 learning: A systematic review and meta-analysis of the experimental evidence. *American Educational Research Journal*, 61(1), 74–107.
- Pan, L., Li, C., Li, J., & Tang, J. (2017). Prerequisite relation learning for concepts in MOOCs. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics* (pp. 1447–1456).
- Piech, C., Bassen, J., Huang, J., Ganguli, S., Sahami, M., Guibas, L., & Sohl-Dickstein, J. (2015). Deep knowledge tracing. In *Advances in Neural Information Processing Systems* 28.
- Prihar, E., et al. (2022). Exploring common trends in online educational experiments. In *Proceedings of the 15th International Conference on Educational Data Mining*.
- Renkl, A. (2014). Toward an instructionally oriented theory of example-based learning. *Cognitive Science*, 38(1), 1–37.
- Rodrigues, L., et al. (2024). Mathematics intelligent tutoring systems with handwritten input: A scoping review. *Education and Information Technologies*, 29, 11183–11209.
- Roediger, H. L., & Karpicke, J. D. (2006). Test-enhanced learning: Taking memory tests improves long-term retention. *Psychological Science*, 17(3), 249–255.
- Rupp, A. A., Templin, J., & Henson, R. A. (2010). *Diagnostic measurement: Theory, methods, and applications*. Guilford Press.
- Sadler, P. M. (1998). Psychometric models of student conceptions in science: Reconciling qualitative studies and distractor-driven assessment instruments. *Journal of Research in Science Teaching*, 35(3), 265–296.
- Sadler, P. M., & Sonnert, G. (2016). Understanding misconceptions: Teaching and learning in middle school physical science. *American Educator*, 40(1), 26–32.
- Salas-Pilco, S. Z., et al. (2024). Artificial intelligence for enhancing special education for K-12: A decade of trends, themes, and global insights (2013–2023). *International Journal of Artificial Intelligence in Education*, 35, 1129–1177.
- Scheines, R., Silver, E., & Goldin, I. (2014). Discovering prerequisite relationships among knowledge components. In *Proceedings of the 7th International Conference on Educational Data Mining*.
- Schmidt, W. H., Wang, H. C., & McKnight, C. C. (2005). Curriculum coherence: An examination of US mathematics and science content standards from an international perspective. *Journal of Curriculum Studies*, 37(5), 525–559.
- Settles, B., & Meeder, B. (2016). A trainable spaced repetition model for language learning. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics* (pp. 1848–1858).
- Shute, V. J. (2008). Focus on formative feedback. *Review of Educational Research*, 78(1), 153–189.
- Siegler, R. S., Duncan, G. J., Davis-Kean, P. E., Duckworth, K., Claessens, A., Engel, M., Susperreguy, M. I., & Chen, M. (2012). Early predictors of high school mathematics achievement. *Psychological Science*, 23(7), 691–697.
- Steenbergen-Hu, S., & Cooper, H. (2013). A meta-analysis of the effectiveness of intelligent tutoring systems on K–12 students' mathematical learning. *Journal of Educational Psychology*, 105(4), 970–987.
- Student Achievement Partners. (2015). *Coherence map*. Achieve the Core. <https://tools.achievethecore.org/coherence-map/>

- Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. *Cognitive Science*, 12(2), 257–285.
- TNTP. (2018). *The opportunity myth: What students can show us about how school is letting them down—and how to fix it*. TNTP.
- TNTP & Zearn. (2021). *Accelerate, don't remediate: New evidence from elementary math classrooms*. TNTP.
- Treagust, D. F. (1988). Development and use of diagnostic tests to evaluate students' misconceptions in science. *International Journal of Science Education*, 10(2), 159–169.
- U.S. Department of Education, Office of Educational Technology. (2023). *Artificial intelligence and the future of teaching and learning: Insights and recommendations*.
- U.S. Department of Education, Privacy Technical Assistance Center. (n.d.). *Protecting student privacy while using online educational services: Requirements and best practices*. <https://studentprivacy.ed.gov/>
- VanLehn, K. (2011). The relative effectiveness of human tutoring, intelligent tutoring systems, and other tutoring systems. *Educational Psychologist*, 46(4), 197–221.
- Vygotsky, L. S. (1978). *Mind in society: The development of higher psychological processes*. Harvard University Press.
- Wang, R. E., Zhang, Q., Robinson, C., Loeb, S., & Demszky, D. (2024a). Bridging the novice-expert gap via models of decision-making: A case study on remediating math mistakes. In *Proceedings of NAACL 2024*.
- Wang, R. E., Ribeiro, A. T., Robinson, C. D., Loeb, S., & Demszky, D. (2024b). *Tutor CoPilot: A human-AI approach for scaling real-time expertise* (EdWorkingPaper 24-1054). Annenberg Institute, Brown University.
- Wiggins, G., & McTighe, J. (2005). *Understanding by design* (expanded 2nd ed.). ASCD.
- Wood, D., Bruner, J. S., & Ross, G. (1976). The role of tutoring in problem solving. *Journal of Child Psychology and Psychiatry*, 17(2), 89–100.
- Zimba, J. (2012). *Wiring diagram of the Common Core State Standards for Mathematics*. Posted at commoncoretools.me.