Inverting the Priority Stack:

Rights-Based AI Tutoring for Special Education Through Structural Design

Travis Gilly
ORCID: 0009-0007-2954-6313
Real Safety AI Foundation
Illinois, United States

Working paper, June 2026 (v1)

Abstract

Current AI-enabled educational technology for special education optimizes for district compliance rather than student learning, treating the child as a data subject and the parent as a notification endpoint, a pattern the critical literature on the datafication of education has documented for over a decade. Teacher in the Loop (TITL) is a patent-pending, design-stage AI tutoring platform built exclusively for special education that inverts this priority structure: the student receives Socratic instruction calibrated to their cognitive profile, the parent controls engagement frequency through a portal grounded in IDEA 2004, the teacher retains pedagogical authority backed by a contractual anti-displacement clause, and the district receives compliance documentation as an automatic byproduct rather than the system's purpose. The paper operationalizes this inversion through six rights-based features, each embedding a specific protection at the structural layer rather than as a configurable option. Drawing on meta-analytic evidence for intelligent tutoring and on research showing that algorithmic classification errors carry heavier consequences for the students least equipped to absorb them, it argues that placing children's rights in the architecture, rather than bolting them onto systems designed for institutional convenience, produces better outcomes for every stakeholder.

Keywords: special education, intelligent tutoring systems, children's rights, structural design, datafication, human oversight, accommodation, IDEA 2004, privacy by design

I. Introduction

The integration of artificial intelligence into special education has followed a pattern familiar across educational technology: systems are designed for the entity that signs the purchase order. In the United States, that entity is the school district. A district does not buy software to be tutored. It buys software to manage risk, to document accommodation delivery, to satisfy auditors, and to reduce the administrative weight of legally mandated services. When a product is built to be sold to a district, the district's needs become the product's center of gravity, and everyone else in the room gets arranged around them. This is the standing critique of educational technology built for institutional buyers rather than for learners (Selwyn, 2022; Williamson & Hogan, 2021).

This produces a recognizable ordering of concern. The district sits at the top, because the district pays. The teacher sits below, valued mainly as the operator who keeps the district's data flowing. The parent sits lower still, treated as someone to be notified after decisions are made. The student, the one person the system exists to serve, sits at the bottom, present in the architecture as a data subject whose interaction is the raw material the compliance machine consumes. This is not figurative. A substantial literature on the datafication of education documents how digital school systems render children as streams of machine readable data, gathered across every level of the system and repurposed for accountability and management rather than for the child (Jarke & Breiter, 2019; Mascheroni, 2018). I call this ordering the priority stack, and in conventional special education technology it runs district first, student last.

The claim of this paper is that the stack is upside down, and that turning it over is a design decision rather than a slogan. Teacher in the Loop (TITL) is a design-stage adaptive tutoring platform built only for students receiving legally mandated accommodations. It is constructed on the inverted stack: the student is first, the parent is a participant rather than a recipient, the teacher holds authority rather than operating a dashboard, and the district receives what it needs as a byproduct of a system whose actual purpose is teaching the child. The contribution here is not an efficacy demonstration, which deployment will have to settle, but an argument about where rights belong in the architecture. The framing follows an emerging agenda that treats children's rights as a first order concern in the design of the digital environments children inhabit rather than as an afterthought (Livingstone & Third, 2017). Rights placed at the structural layer behave differently from rights bolted on as features, and the difference shows up in what each stakeholder can and cannot do once the system is running.

The paper proceeds as follows. Section II describes the conventional priority stack and the harms that follow from it. Section III states inversion as a design commitment rather than a marketing posture. Section IV draws the distinction between structure and bolt-on, which is the load-bearing distinction of the paper. Section V walks the inverted stack stakeholder by stakeholder and ties each position to a concrete feature of the platform. Section VI describes the design origins, because the architecture came from a particular vantage and that vantage matters to the argument. Section VII explains why inversion improves outcomes for every stakeholder, including the district that conventional design tries hardest to please. Section VIII states the limitations honestly, including the fact that the platform is at design stage. Section IX concludes.

II. The Conventional Priority Stack

Special education technology in the United States operates inside a legal frame defined by the Individuals with Disabilities Education Improvement Act of 2004 (IDEA), which guarantees a free appropriate public education and requires that services be specified in an individualized education program. The statute creates real obligations and real liability. A district that fails to deliver the accommodations written into a child's program is exposed, and that exposure is the lever every special education vendor pulls. The pitch is rarely that the software teaches better. The pitch is that the software protects the district.

Once protection is the selling point, the architecture organizes itself around documentation. The system's first job becomes producing a record that the district can show to an auditor or a hearing officer, and the student's learning becomes the occasion for generating that record rather than the goal the record serves. The child is rendered, in the data model, as a stream of events to be logged. The parent is rendered as an address to which notices are sent. The teacher is rendered as a set of fields to be completed. None of these renderings is hostile. They are simply what happens when the entity that pays is the entity the system is built to satisfy.

The harms are not dramatic and they are not accidental. They are the ordinary consequence of the ordering. Empirical work on datafied schooling gives each rendering a face. Teachers report that the pressure to produce data can be experienced as detrimental to the work of building secure relationships with the children the data is about (Jarke & Breiter, 2019). Behavior tracking platforms that push daily updates to parents on the platform's schedule have been shown to reconfigure the relationships among teacher, parent, and child around the data flow itself rather than around the child (Jarke & Breiter, 2019). And the datafied child becomes an object of routine monitoring normalized across institutional and family life, a condition the critical literature on childhood datafication has tracked for the better part of a decade (Mascheroni, 2018). A platform that treats the child as a data subject collects more than instruction requires, because more data is a better record. A platform that treats the parent as a notification endpoint pushes information outward on the platform's schedule rather than giving the parent any control over the flow, because control by the parent is friction in the documentation pipeline. A platform that treats the teacher as an operator routes authority away from the person with the most context about the child, because the operator's judgment is a variable the compliance machine would rather not depend on, a dynamic that turns the teacher into a kind of surveillant consumer of the platform rather than its pedagogical authority (Kumar et al., 2019). Each of these is defensible in isolation and indefensible in aggregate, and the aggregate is the system the field currently ships.

III. Inversion as a Design Commitment

To invert the stack is to decide, before any feature is specified, that the ordering will run student, parent, teacher, district, and to hold that ordering when features compete for priority. This is a commitment that has to survive contact with engineering tradeoffs, because every real system forces choices about what is on by default, what is optional, who can override whom, and what the system does when two interests collide. A platform that says it puts students first and then resolves every tradeoff in the district's favor has not inverted anything. It has printed a new slogan over the old architecture.

Inversion is therefore testable by inspection. One does not need deployment data to ask, of any given feature, whose interest it serves when interests diverge. If a sensing feature collects data the child's instruction does not require, the inverted stack turns it off by default, because the student's privacy outranks the district's appetite for a richer record. If a notification feature floods a parent on the platform's schedule, the inverted stack hands the schedule to the parent, because the parent's participation outranks the pipeline's convenience. If a classification feature would let the system route a child without a human reading the context, the inverted stack routes the recommendation to the teacher, because the teacher's authority outranks the efficiency of an automated decision. The ordering is not a value statement appended to the documentation. It is the rule that decides the tradeoffs, and it is visible in the defaults.

IV. Structure Versus Bolt-On: Where Rights Live

The central distinction of this paper is between a right that lives in the structure of a system and a right that is bolted on as a feature. The difference is what happens under pressure. A bolted-on right is a setting, and a setting can be changed, defaulted off, buried, or quietly dropped in the next release when it becomes inconvenient. A structural right is a property of how the system is built, such that removing it would require rebuilding the system rather than flipping a toggle. Privacy stated in a policy document is a bolt-on. Privacy enforced by processing student data only in tokenized form, with a locally deployable option that keeps inference on a district-controlled server, is structural, in the sense that privacy becomes a property of the build rather than a promise layered on top of it (Cavoukian, 2009), because the data the policy promises to protect is never assembled in the first place.

This distinction matters because special education technology is full of bolted-on rights. A platform will advertise human oversight, parental control, and data minimization, and then implement each as a feature that the operating institution can configure away. Recent work on inclusive classroom technology finds that oversight stated in documentation deteriorates into language when it is not enforced by the architecture, and that the interpretive work of the human educator is the variable that decides whether an AI system produces inclusive outcomes or categorization effects (Pasqualotto & Marras, 2026). The lesson is that the claim and the architecture can diverge, and that a reviewer who reads only the claim will not see the divergence. The cure is to build the right into the structure, so that the claim and the architecture are the same object.

This structural turn has a direct precedent in the privacy literature. Hull (2021), in his account of what he calls the death of the data subject, argues that the rights bearing data subject is being disempowered in a way that piecemeal vindication of individual rights cannot repair, because the disempowerment is produced by the power asymmetry between the individual and the entity holding the data, and the remedy must therefore be structural rather than a matter of asserting rights one at a time. The point transfers cleanly to a child inside a special education platform. The asymmetry between a school district and a child is close to total, and a right the child holds only as a setting the district is free to configure is not a right the child can rely on. Placing the right in the structure, where exercising it does not depend on the more powerful party choosing to honor it, is what makes the right real.

TITL applies this cure across its feature set. Teacher authority is structural because the AI operates as a suggestion layer rather than a decision layer, and the teacher receives full context on every recommendation and retains the power to override any system-generated classification. Labor protection is structural because a contractual anti-displacement clause commits any implementing institution to not reduce teaching, aide, or special education personnel as a result of efficiency gains, which makes the platform's central premise enforceable rather than aspirational. Privacy is structural because non-essential data is never collected by default. In each case the right is not a setting the district can change. It is a property of the build.

V. The Stack Inverted, Stakeholder by Stakeholder

A. *The Student, First*

Placing the student first means the system's primary obligation is instruction matched to the individual child, and that the accommodations written into the child's individualized education program are treated as structurally required behavior rather than optional configuration. A small set of commitments are non-negotiable for exactly this reason. Socratic only instruction, the hybrid privacy layer described below, and core accommodation delivery are active for every child regardless of configuration, because they protect the child's rights rather than enrich the record. A child does not opt in to being taught.

The instructional core is a Socratic engine. The tutor does not hand over answers. It proceeds by questioning, analogy, counterexample, and scaffolded hints calibrated to the student's current understanding, which protects the child's right to genuine learning rather than answer delivery. Its design principle is to hold the student in the zone of proximal development, the space where a task is hard enough to demand real effort but not so hard that effort collapses, the productive difficulty the learning sciences associate with durable gains rather than the unproductive failure that follows when challenge outruns support (Kapur, 2016). Difficulty adjusts continuously from non-invasive behavioral signals: time on a problem relative to the student's own baseline, error rate trajectory, response latency, and help seeking patterns. There are no cameras and no biometric sensors. When the signals show a student crossing from productive struggle into shutdown, a three-tier path routes them to progressively more human support, from AI hint reframing, to peer matching, to direct teacher notification carrying the full context of what the student has attempted.

The differentiation is profile specific rather than generic. A student with autism receives structural consistency in how material is presented and how interactions unfold; a student with ADHD receives varied delivery that sustains attention through novelty. This adjustment runs automatically from the child's program profile and does not require real time teacher configuration. The claim is not that the platform has proven these gains for this population. The evidence that tutoring systems of this kind produce real learning is strong, but it is drawn largely from general education. Two meta-analyses anchor it. Kulik and Fletcher (2016), across fifty controlled evaluations, found a median gain of roughly two thirds of a standard deviation, moving a median student from the fiftieth to the seventy fifth percentile.

Ma, Adesope, Nesbit, and Liu (2014), pooling 107 effect sizes across more than fourteen thousand learners, found intelligent tutoring statistically indistinguishable from one to one human tutoring. Transfer to students with documented learning needs is a hypothesis the platform is built to test, not a result in hand. Putting the student first does not mean overstating what is known on the student's behalf.

A second student facing feature is invisible accommodation delivery. When a child logs in, the program's accommodations, extended time, modified presentation, alternative response formats, and reading supports, activate without any visible marker, so that an accommodated session is indistinguishable from a standard one to the student and to peers. This addresses a documented barrier. Students who know an accommodation is active often decline it because of stigma, an effect that intensifies in adolescence (Toutain, 2019). Making the accommodation invisible removes the social cost of receiving legally mandated support. Each delivery is logged with a timestamp, which incidentally produces a continuous verification record where current practice relies on teacher attestation.

B. The Parent, as Participant Rather Than Recipient

In the conventional stack the parent is an endpoint, a destination for notices generated on the platform's schedule. IDEA 2004 says otherwise. Its provision for alternative methods of progress communication with parental consent makes the parent a participant rather than a recipient (20 U.S.C. § 1414(d)). The inverted stack builds the parent a portal grounded in that standing, showing daily instructional activity, accommodation delivery verification, and progress toward program goals, with the parent, not the platform, setting the frequency, whether daily, weekly, or summarized. The control sits with the person the statute names as a participant. In a preliminary consultation with twenty-five families in a special education program, every family chose daily updates once the choice was actually offered, which is itself a comment on how poorly the periodic notification model serves the people it claims to inform.

The parent's authority also governs the line between essential and non-essential capability. The non-negotiable baseline aside, the majority of the platform's features are explicitly opt-in or opt-out, the choice belonging to the parent and revocable at any time. During onboarding the parent receives a plain language account of what each feature would do for their specific child within the program. A parent uneasy about a capability such as cognitive profile modeling or process capture can decline it without losing the rest of the platform. This is not permissiveness. It is the operationalization of parental authority, on the premise that the parent, not the district, is the appropriate decision maker for a child's engagement with AI.

C. The Teacher, as Authority Rather Than Operator

The platform's name states its governance principle. The teacher is in the loop not as a courtesy but as the holder of pedagogical authority at every level, with the AI acting as an assistive layer that extends the teacher's judgment rather than substituting for it. Recent work on AI in inclusive classrooms finds that this is the variable that decides outcomes. Adoption is legitimate only when the technology is aligned with professional control, governance clarity, and ethical accountability, and sustainable integration depends less on technological expansion than on teachers' informed professional judgment (Pasqualotto & Marras, 2026). The teacher dashboard provides real time visibility into student activity, struggle classification, intervention history, and accommodation delivery, and the teacher retains the authority to override any classification the system produces. The point of the dashboard is not to monitor the teacher. It is to give the teacher the context the system has gathered, so that the human with the most knowledge of the child makes the calls that matter.

This matters because misclassification here is not a minor inconvenience. Adaptive learning systems that use machine learning to infer mastery, engagement, and affect are not equally effective for all groups of students, and the distinctive danger of automated classification is the scale at which a wrong inference can propagate (Belitz et al., 2022). A child mislabeled by an autonomous system can be routed wrongly at a speed no individual teacher could match. The mitigation in TITL is architectural rather than procedural. The AI proposes and the teacher disposes, the teacher sees the full context behind each recommendation, and no system-generated label routes a child without a human positioned to correct it. This does not eliminate error. It confines error to the zone where teacher judgment is available as a correction, which is the most an honest design can promise and considerably more than a system that lets classifications act on their own.

A second teacher facing feature is process capture. The writing environment records the student's revision process, the drafts, the edits, and the time between them, not as surveillance but as the student's own evidence of authorship, which reframes academic integrity from a policing function into a demonstration that the work is genuinely the student's. The teacher's authority is reinforced by a contractual anti-displacement clause: any institution that reduces teaching, aide, or special education staff as a result of efficiency gains from the platform loses access to it. This is not a feature in the technical sense. It is a right written into the terms of use, protecting the workforce the system was built to support rather than replace.

D. The District, Served by Byproduct

The inverted stack does not ignore the district. It serves the district last, and it serves it well, by producing the documentation the district needs as a byproduct of teaching rather than as the system's purpose. The audit trail that TITL generates, accommodation delivery logging, suggestion tracking, and a record of teacher interventions, constitutes a compliance posture that exceeds current special education documentation standards. The district gets a stronger record than a compliance-first product would give it, and it gets that record without the system having been organized around producing it.

This is the move that makes inversion practical rather than merely principled. A platform that served the student first and left the district with nothing would not be bought, and a tool that is not bought helps no child. The inverted stack avoids that failure by recognizing that good teaching, fully documented, produces exactly the evidence the district requires. Compliance is treated throughout as a byproduct rather than a product category. The platform does not market itself as a compliance tool. It markets itself as a pedagogical tool whose ordinary operation happens to generate the evidence districts need, which is a different posture from the compliance-first competitors and a better one for the child, because the child's learning is the thing being optimized and the record is the residue.

VI. Design Origins

This architecture did not come from a market analysis. It came from the experience of a neurodivergent parent navigating the individualized education program process for a neurodivergent child, and watching a system that was supposed to serve that child spend most of its energy serving itself. The vantage matters to the argument because the priority stack is most visible from the bottom. A district sees a compliance burden and a tool that lightens it. A parent at an individualized education program meeting sees a child being processed and a set of decisions arriving as notifications. The inverted stack is what the architecture looks like when it is designed from the parent's chair rather than the procurement office, and the features described above are the concrete form of that perspective.

I raise this not as biography but as method. Selection of whose interests anchor a system is the most consequential design decision in educational technology, and it is usually made implicitly by the question of who pays. Making it explicitly, and anchoring it on the student, is the whole of the design philosophy. The lived vantage is the reason the anchoring was obvious to me when it is apparently not obvious to a field that keeps shipping the stack upside down.

VII. Why Inversion Improves Outcomes for Everyone

It would be easy to read inversion as a transfer, a redistribution that helps the student at the district's expense. It is not. The argument of this section is that the inverted stack produces better outcomes at every level of the stack, including the level conventional design works hardest to please.

The student gains instruction that is actually optimized for learning rather than for record generation, and privacy that is a property of the build rather than a promise in a policy. The parent gains a participatory role the statute already grants and the conventional stack already ignores, with control over the flow of information and over the data the system collects. The teacher gains authority, context, and protection from displacement, which together preserve the conditions under which good teaching can happen at all. And the district gains a documentation posture stronger than the compliance-first tools provide, produced as a byproduct of a system that is genuinely teaching the children whose records it generates. The district that buys a compliance tool gets a record. The district that buys the inverted tool gets a record and a learning gain, and the record is better because it documents real instruction rather than instruction shaped to be documented.

This is why inversion is a design argument and not a moral concession. The conventional stack is not merely unjust. It is inefficient at the thing it claims to optimize, because a system built to produce records produces worse learning than a system built to teach, and worse records too, since the best record is the residue of real teaching. Putting the student first is the design that serves the district best. That the field has not seen this is a fact about who it builds for, not a fact about what works.

VIII. Limitations

The most significant limitation of this paper is that TITL is at design stage and has no deployment data. Every claim about outcomes is therefore an argument about architecture rather than a report of measured effect, and the efficacy of the inverted stack in practice is an empirical question that a controlled deployment in a special education context will have to answer, with accommodation delivery verification, measures of teacher agency, and student outcome measures as the endpoints that matter.

Two further limitations deserve naming. First, the evidence that adaptive AI tutoring produces strong learning gains comes from general education settings, and its transfer to students with documented learning needs is a hypothesis rather than an established extrapolation. The mechanisms are plausible and the design principles translate, but the population is different and the claim should be held at the level of hypothesis until the target population is studied directly. Second, the strategy of adapting instructional content differentially across heterogeneous neurodivergent profiles is the platform's least settled commitment, and it should be presented as a theoretically motivated design to be tested rather than a validated finding. A platform that presents a hypothesis as a result invites the reviewer reaction that has damaged other AI education efforts, and the inverted stack has no need to overclaim, because its argument is structural and does not rest on outcomes it has not yet measured.

IX. Conclusion

Educational technology for special education is built upside down. It is built for the entity that pays, and so it ranks the district first and the student last, with the parent and the teacher arranged in between as a recipient and an operator. The harms that follow are not failures of intention. They are the ordinary consequences of the ordering, and they will persist for as long as the ordering does. Teacher in the Loop is an attempt to turn the stack over and to hold it over under engineering pressure, by building the student's instruction, the parent's participation, and the teacher's authority into the structure of the system rather than bolting them on as features the operating institution can configure away. The district is served last and served well, by a record that is the byproduct of real teaching. Whether the platform teaches as well as the architecture intends is a question for deployment. Whether the architecture is pointed at the right person is a question of design, and on that question the inverted stack gives an answer the field has been avoiding: the system should be built for the child it exists to serve.

Acknowledgments

The author discloses the use of speech-to-text software (Wispr Flow) and Anthropic's Claude as assistive technology for a diagnosed neurodevelopmental condition. All intellectual contributions, research questions, theoretical arguments, methodological decisions, and conclusions are solely the author's own.

Correspondence concerning this article should be addressed to Travis Gilly, Real Safety AI Foundation, Illinois, United States. Email: `t.gilly@ai-literacy-labs.org`.

References

Belitz, C., Ocumpaugh, J., Ritter, S., Baker, R. S., Fancsali, S. E., & Bosch, N. (2022). Constructing categories: Moving beyond protected classes in algorithmic fairness. *Journal of the Association for Information Science and Technology*, 74(6), 663–668. https://doi.org/10.1002/asi.24643

Cavoukian, A. (2009). *Privacy by design: The 7 foundational principles*. Information and Privacy Commissioner of Ontario.

Hull, G. (2021). The death of the data subject. *Law, Culture and the Humanities*, 20(3), 527–547. https://doi.org/10.1177/17438721211049376

Individuals with Disabilities Education Improvement Act of 2004, Pub. L. No. 108-446, 118 Stat. 2647, 20 U.S.C. § 1414(d).

Jarke, J., & Breiter, A. (2019). Editorial: The datafication of education. *Learning, Media and Technology*, 44(1), 1–6. https://doi.org/10.1080/17439884.2019.1573833

Kapur, M. (2016). Examining productive failure, productive success, unproductive failure, and unproductive success in learning. *Educational Psychologist*, 51(2), 289–299. https://doi.org/10.1080/00461520.2016.1155457

Kulik, J. A., & Fletcher, J. D. (2016). Effectiveness of intelligent tutoring systems: A meta-analytic review. *Review of Educational Research*, 86(1), 42–78. https://doi.org/10.3102/0034654315581420

Kumar, P., Vitak, J., Chetty, M., & Clegg, T. L. (2019). The platformization of the classroom: Teachers as surveillant consumers. *Surveillance & Society*, 17(1/2), 145–152. https://doi.org/10.24908/ss.v17i1/2.12926

Livingstone, S., & Third, A. (2017). Children and young people's rights in the digital age: An emerging agenda. *New Media & Society*, 19(5), 657–670. https://doi.org/10.1177/1461444816686318

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. https://doi.org/10.1037/a0037123

Mascheroni, G. (2018). Datafied childhoods: Contextualising datafication in everyday life. *Current Sociology*, 68(6), 798–813. https://doi.org/10.1177/0011392118807534

Pasqualotto, A., & Marras, A. (2026). Negotiating assistive technologies and AI in inclusive education: Professional agency in neurodevelopmental contexts. *Frontiers in Child and Adolescent Psychiatry*, 5. https://doi.org/10.3389/frcha.2026.1820276

Selwyn, N. (2022). *Education and technology: Key issues and debates* (3rd ed.). Bloomsbury Academic.

Toutain, C. (2019). Barriers to accommodations for students with disabilities in higher education: A literature review. *Journal of Postsecondary Education and Disability*, 32(3), 297–310.

Williamson, B., & Hogan, A. (2021). *Pandemic privatisation in higher education: Edtech and university reform*. Education International Research.