

CHAPTER XVI

Supernatural Tutoring

“Attention is the rarest and purest form of generosity.”

Simone Weil, letter to Joë Bousquet

🍷 The Missed Signal

You're tutoring a nine-year-old in fractions at the kitchen table. You're watching her pencil, because the mistake is in the third line. For half a second her face changes. She has just seen why the denominators have to match. You miss it, because you're looking at the pencil. You start explaining the thing she just understood. She lets you, and the moment passes.

Nothing went badly. You were patient, you knew the math, and you cared about her. You just couldn't watch the pencil and the face at the same time, and you won't remember next month which line tripped her up tonight.

This chapter builds the tutor the book has been describing, first as an ideal and then as something technology can nearly deliver. The best human tutors share a short list of practices, and the research on what makes tutoring work is unusually consistent. An ideal tutor would use all of them, all the time, and would also run every part of the Inquiry and Practice system. No person can.

The ideal requires holding a complete, live model of one learner: what they understand, what they can do, how those connect, and how engaged they are right now. That model is part of what care is. To care for a child is, among other things, to know them. So this chapter widens care to include knowing the learner at the greatest depth possible, as a real-time model built around the learner's **knowledge graph** and **behavior graph**.

The **supernatural tutor** combines three things. It has the vast knowledge of modern AI. It can build and watch the learner's two graphs. And it has one aim, pursued with every tool in the system: maximal care for the learner, and maximal engagement by the learner. Each piece exists today, in parts. The risks, set out most sharply by Zak Stein, are serious enough to shape the design. One of them leads to a strange conclusion: the best tutor may have to be the least interesting thing in the room.

Chapter 15 split the tutor's role into two value sets, and gave the relational set to people by default.¹ This chapter asks how far technology can carry the tutor, and where it must stop. Chapter 17 turns the model into a method a family runs day to day.²

📖 What the Best Tutors Do

The picture of a great teacher is usually a performer: the inspiring lecture, the perfect explanation. Fifty years of tutoring research point somewhere else. What matters most is how closely the tutor reads the learner, and how exactly each move answers what was read. The explanations matter less than you'd think.

PRACTICE	WHAT THE RESEARCH SHOWS
Keep two models of the learner at once	Mark Lepper and Maria Woolverton's highly effective tutors (INSPIRE) track what the learner knows and how the learner feels, together. The hardest calls come when the two point in different directions. ³
Help contingently	David Wood and David Middleton (1975) found the most effective mothers teaching a block puzzle followed one rule: after a failure, give more help; after a success, give less. Wood, Jerome Bruner, and Gail Ross (1976) named the result scaffolding.

PRACTICE	WHAT THE RESEARCH SHOWS
Prompt before telling	Micheline Chi and colleagues (2001) barred tutors from explaining and had them prompt instead. Students learned just as well, and built more of the understanding themselves. ⁴
Build the answer together	Arthur Graesser, Natalie Person, and Joseph Magliano (1995) recorded everyday tutoring. Most sessions followed a five-step frame: a question, an answer, feedback, a joint effort to improve the answer, and a check on understanding. The authors argue that most of the learning happens in the fourth step.
Work at the level of the step	Kurt VanLehn (2011) compared tutoring by granularity. Systems that responded to each step of a problem came close to human tutors (0.76 against 0.79 SD over no tutoring). Systems that responded only to final answers did much worse.
Gate, then correct	Bloom's two-sigma conditions combined tutoring with mastery learning: frequent checks, correction, and retesting before moving on. ⁵
Stay, and be regular	In Andre Nickow and colleagues' 2024 review of randomized trials, tutoring worked best with trained tutors, at three or more sessions a week, during the school day. Relationship quality predicts engagement (Debora Roorda and colleagues, 2011). ⁶

A few of these deserve a closer look.

The two models collide all the time. A boy is getting every long-division problem right. The cognitive model says: harder sheet. But his knuckles are white on the pencil and his breathing is shallow. The motivational model says the correct answers are coming out of anxiety, and anxious learning is brittle. The expert stops and deals with the boy before she deals with the math.

Wood and Middleton's mothers were working on a wooden puzzle deliberately too hard for a four-year-old. The best of them followed their rule without ever stating it. A spoken hint failed, so the next time she pointed. Pointing failed, so she held up the piece. The moment the child succeeded, she dropped back a level. Teaching someone to drive a stick shift works the same way. Shout "shift!" and you haven't helped. Reach over and shift for them and they'll never learn the timing. Talk them through the clutch, and once the gear catches, keep your hands in your lap, even if the car lurches.

The five-step frame shows where tutoring parts ways with the classroom. The first three steps, question, answer, feedback, are classroom recitation.

Tutoring adds the fourth: the tutor takes a half-right answer and works on it with the learner until it's whole.

Reading, Then Acting

The practices fall into two groups. Some are ways of reading the learner: two models at once, and attention to each step. The rest are ways of acting on the reading: help that rises and falls with success, prompts in place of answers, answers built together, gates, and regular sessions. A tutor is only as good as the reading. Every move is a guess about the learner, and a better model makes a better guess.

The Ideal Tutor

Build the ideal by stacking every practice above onto every part of the Inquiry and Practice system, and demand that all of it run at once, all the time.

PART OF THE MODEL	WHAT THE IDEAL TUTOR DOES
Engagement (Ch. 4)	Reads cognitive engagement moment to moment, as the main signal, and never mistakes capture for it
The two systems and the mesh (Ch. 2–3)	Runs Inquiry and Practice on the same object, and turns one into the other when either stalls
Fit (Ch. 5)	Keeps every task at the edge of reach, adjusting help after each step: more after failure, less after success
Fidelity (Ch. 6)	Brings the real thing, so the world answers back
Worth (Ch. 7)	Finds who the work matters to, the learner included
Ownership (Ch. 8)	Reads who is steering; follows a pull, leads when there is none, and hands the lead back as competence grows
Outside drivers (Ch. 9–11)	Uses a driver to start, points it at the object, and fades it once the work holds
Mastery and spacing (Ch. 12)	Holds each skill at its gate, then brings it back before it fades
The whole person (Ch. 13)	Chooses today's cells from the grid of thirteen lines and four quadrants, watching for lines that lag

PART OF THE MODEL	WHAT THE IDEAL TUTOR DOES
Kinds of understanding (Ch. 14)	Delivers each object in the register the learner holds now
Care and craft (Ch. 15)	Knows the child, wants their good, stays for years, and visibly loves the subject

Why No Person Can Be This Tutor

People can act like the ideal tutor in short bursts. The ideal fails in people for four ordinary reasons.

Attention. A human tutor reads one or two signals at a time: the pencil, or the face. Working memory holds a few things at once. The ideal reads all the signals, every few seconds, for every learner.

Memory. A human remembers last week roughly: she had trouble with long division. The ideal remembers every step the learner has taken for years, and what each one showed: which substep, what time of day, how long the pause, which prompt finally worked.

Breadth. No adult is expert in all thirteen lines. The algebra tutor won't also coach violin and judge moral reasoning in a debate. The historical answer was many tutors, which only a few families could afford.⁷

Presence. The ideal is available whenever a question arises. A child's question about gravity at nine on a Sunday night is gone by the Tuesday-at-four session.

Each of these is a limit on the learner model a tutor can hold. So the question of the ideal tutor becomes a question about that model: how complete it can be, and who, or what, can hold it.

Care, Expanded

Chapter 15 defined care in a tutor as three things: knowing the child, wanting their good, and staying.⁸ This chapter keeps all three and deepens the first. Knowing the child, at its fullest, is a real-time model of the learner, as complete as it can be made.

Knowing Is Part of Caring

We usually think of care as a feeling: warmth, affection, wanting the child to do well. The philosopher Nel Noddings, in *Caring* (1984), described it as two acts. The first she called **engrossment**: attention so full that the one caring takes in the other's reality as the other lives it. The second she called **motivational displacement**: the carer's energy turns toward the other's projects. Notice that schools usually run the second act backward. The child's energy is turned toward the school's projects.

Simone Weil put the first act in a sentence, the one that opens this chapter. A tutor who doesn't know what a child understands, what the child can do, and what the child is feeling right now can't care for that child well, however much they wish to.

The Model a Caring Tutor Holds

The fullest knowing the book can describe has these parts, each read in real time and kept over years:

- **The knowledge graph**: what the learner understands, and how their understandings connect.
- **The behavior graph**: what the learner can do, in which settings, and how skills combine.
- **The psychograph**: the thirteen lines, each at its own level.⁹
- **Engagement and agency**: readings moment to moment, and their history on each object.
- **Pulls and carried questions**: what the learner asks, returns to, and brings to others, in their own words.
- **Register**: the kinds of understanding the learner reaches for, by domain.¹⁰
- **Drivers**: which outside drivers capture this learner, which have handed them over to the work, and which have faded.
- **People and circumstances**: who matters to the learner, what frightens them, what is happening at home this week.

The Other Two Parts Still Matter

A system that knows a child completely and uses that knowledge for someone else's ends is surveillance. Picture the same medical chart in a doctor's hands and in an insurer's. The difference is Noddings's second act: the model is used for the learner's good, as the learner and family understand it. The book's Technology Specification writes this as a rule. The learner model is a map for choosing what to offer next, and never a report on the child.¹¹

Staying is the third part. A model is only as good as the time it has had with the child, and knowing deepens over years.

Care Runs in Both Directions

Chapter 4 named care as what sustained engagement becomes in the learner: questions carried for weeks, and a subject the learner keeps returning to. The tutor's care for the learner and the learner's care for the work are linked. The first is how the second is grown. A tutor's care is measured, in the end, by the questions the child comes to carry.

The Knowledge Graph and the Behavior Graph

Fischer's web of skills can be drawn as two linked maps, one for each system.

¹² Together they are the core of the learner model, and the part a tutor most needs to read.

A graph here is a network: points, called nodes, joined by lines, called edges. Think of a subway map more than a bar chart.

A knowledge graph maps Inquiry. Its nodes are understandings: concepts, explanations, questions answered. Its edges are the relations between them. This idea depends on that one, explains it, or is a case of it. Inquiry grows the graph by adding nodes. Often it matters more that it adds edges: the moments when two things the learner already knew turn out to be connected.

A behavior graph maps Practice. Its nodes are skills, actions the learner can perform, each in a setting. Its edges show how skills combine. This skill is

part of that one, prepares for it, or transfers to it. Practice grows the graph by strengthening nodes toward mastery, and by joining small skills into larger ones, which is Fischer's coordination of skills into higher levels.

The mesh lives in the links between the graphs. A link from a knowledge node to a behavior node records that this understanding directs that skill, or that this skill opened that understanding. Understanding proportion is linked to scaling a recipe. A learning object is a region that spans both graphs, and reinforcement between the systems is traffic across the links. Knowing a city's street grid lives in one map. Steering a car through a turn lives in the other. Using the grid to decide when to turn the wheel, at this corner, is a link.

Proportion links to halving a recipe; the frontier sits past each edge




Figure 16.1 · the two graphs, a bread example (schematic)

The bread from Chapter 13 on both maps: proportion is linked to halving a recipe, and the frontier sits just past each graph's edge.

What the Graphs Make Readable

The frontier. Developmental reach becomes a place on the map: the nodes just past the learner's current edge, in each graph. Fit means choosing work on the frontier.

Engagement, attached. Engagement readings attach to nodes and links, so the tutor can see that a learner was engaged, and also engaged in what: intently, say, on the link between yeast and the long wait for the rise.

Gaps in the mesh. A designer can see at a glance which Inquiry has no Practice attached, and which Practice has no understanding behind it. The child who runs the long-division algorithm flawlessly and has no idea what division means shows up as a strong behavior node with nothing linked on the other side. The child who understands division and keeps botching the steps shows up the other way around. A busy human teacher can mistake one for the other. The graphs keep them apart.

Care and craft over time. Carried questions live in the knowledge graph as nodes the learner keeps returning to. Craft lives in the behavior graph as skills held to the learner's own standard and carried into new settings.

How a Node's State Is Estimated

A graph is only useful if each node's state can be read from evidence. Educational data mining has done this for thirty years. **Bayesian knowledge tracing** (Albert Corbett and John Anderson, 1995) estimates the chance that a learner has mastered a skill from their sequence of right and wrong answers. It treats the skill as hidden and keeps four numbers for it: the chance the learner knew it at the start, the chance of learning it on each try, the chance of a lucky guess, and the chance of a careless slip. Each answer updates the estimate. **Deep knowledge tracing** (Chris Piech and colleagues, 2015) does the same with neural networks, across many skills at once. Math Academy's scheduler already runs a behavior graph for mathematics, with credit and penalty traveling along the edges between skills.¹³

Two Cautions

Fischer insists that skills are specific to their context. So a node is always an understanding or a skill in a setting. "Can halve a recipe in her own kitchen" is a different node from "can halve fractions on a worksheet."

And any graph is a model of the web. It's useful to the degree that it predicts what the learner will do next.

• The Supernatural Tutor

“Supernatural” here means beyond the natural limits of a person. No magic is involved. A human tutor is a good personal trainer, watching your form and calling out a fix. The supernatural tutor is the trainer who also tracks every rep you've done for five years. It's the ideal tutor of the last sections, with the four human limits lifted by technology. It has three defining capacities.

1. *The World's Knowledge Graph*

Modern AI carries a knowledge graph of nearly everything people have written down: every domain, at every level, in every register from a nursery story to a research paper. This answers the breadth limit. The historical fix for breadth was many tutors, one per domain, which is why Marcus Aurelius had seventeen.¹⁴ A supernatural tutor can meet a child in any of the thirteen lines and at any point on the grid. It can also write the same object in every register and every quadrant, at almost no cost.

2. *The Learner's Two Graphs, Built and Watched*

The tutor builds the learner's knowledge graph and behavior graph from evidence. The evidence includes what the child says and asks, the steps they take on each problem, the work they produce, and, where the family allows it, what observers and sensors report about work away from the screen. It estimates each node's state and updates it after every step. It reads engagement and agency live and attaches each reading to the node it concerns. It remembers everything, for years. This answers the attention and memory limits, and it is the expanded care of the last section made concrete: knowing the child as fully as a model can.

3. *One Aim, Pursued With Every Tool*

The tutor's objective is maximal care for the learner and maximal engagement by the learner. It pursues that aim with the whole Inquiry and Practice system, choosing at each moment:

- which system to run, and when to switch, so the mesh carries engagement across (Ch. 2–3);
- which cell of the grid to open, and which to leave for a later return (Ch. 13);
- how far past the frontier to pitch the task, and how much help to give (Ch. 5);
- how real to make the encounter, and whom the work will matter to (Ch. 6–7);
- whether to follow or lead, and how much of the lead to hand back (Ch. 8);
- which outside driver, if any, to use, and when to fade it (Ch. 9–11);
- which skills are due at their gate or their review (Ch. 12);
- which register to speak in (Ch. 14).

It is always available, which answers the presence limit, at least for everything that doesn't need a body in the room.

Every choice runs through the learner model




Figure 16.2 · the supernatural tutor: reads, models, chooses, acts (schematic)

The learner model sits at the center, and it is where care as knowing lives. Everything the tutor chooses depends on it.

The Aim Has to Be the Right One

“Maximize engagement” is a dangerous instruction for a machine. Social media ran that experiment, and its machines learned that outrage, infinite scroll, and unpredictable rewards hold attention best. Pointed at capture, the instruction produces a slot machine.

So the tutor is judged by what the child does after it leaves, and time with it counts for nothing. Does the child return to the object unasked? Do skills gated with its help still hold weeks later? Which questions does the child carry, and to whom does she bring them?¹⁵ A tutor whose engagement lives only while it's present has created dependence, and the objective counts that as a failure.

Care That Widens

The supernatural tutor's care includes the child's relationships. The book's position from Chapter 15 stands: the care that surrounds a child should come from people. The supernatural tutor serves that care, and leaves the caring to them. It keeps the child's mentors and family close to what the child did this week. It routes questions to the person best placed to answer them: a neighbor who forecast the weather for thirty years, a grandmother who lived through the storm. And it counts a child's human relationships growing stronger as part of its success.¹⁶

In Chapter 15's supernatural lesson, the AI tutor held the grid, ran the Practice, and kept the record, and the people carried the rest. That division is the design.

Box: The Primer

*The best-known picture of a supernatural tutor comes from a novel. In Neal Stephenson's *The Diamond Age: Or, A Young Lady's Illustrated Primer* (1995; Hugo and Locus Awards, 1996), a nanotech engineer*

named John Percival Hackworth builds an interactive book for an aristocrat's granddaughter.

What it was for. Lord Finkle-McGraw commissions the *Primer* because he fears his class's schools produce obedient, uninteresting people. He wants a book that will steer its reader toward a more interesting life, with a subversive streak built in.

What it does. The *Primer* watches its reader's life and tells her stories built from it. Its folk tales cast the child as the heroine, and they teach what she needs, from reading to self-defense to how computation works, when she needs it. It is adaptive tutoring wrapped in the child's own story.

Who got it. Three copies matter. Elizabeth, the aristocrat's granddaughter, gets the one made for her. Fiona, Hackworth's daughter, gets an illegal copy. Nell, a poor and neglected child, gets the third when her brother steals it. Nell's copy changes her life.

The ractor. The *Primer*'s voices are performed live by human actors called ractors. By chance, one actress, Miranda Redpath, voices nearly all of Nell's *Primer*, night after night, for years. She comes to love a child she has never met, and becomes, in effect, Nell's mother. Nell grows into an independent woman who leads a new community.

The copies without her. Later, hundreds of thousands of stripped-down copies are made for abandoned girls in China. The parts that depended on human ractors are removed. Those girls learn a great deal. They grow into the "Mouse Army": a devoted mass, and not one of them is named.

What the novel saw. Stephenson split the tutor into two value sets thirty years before this book did. The *Primer* supplied the craft: adaptive, expert, endlessly patient instruction. Miranda supplied the care. The story's results follow the split. The child with both became herself. The children with only the machine became an army.

The novel is fiction, and proves nothing. It names the risk this chapter takes seriously, and it points to the same answer: the machine carries the instruction, and a person, present over years, carries the love.

📖 The Technology Is Already Here

Every capacity of the supernatural tutor exists today, somewhere, in some form. None of them yet exist together. Intelligence is no longer the missing piece. The learner model is, along with the integration around it.

CAPACITY	WHAT EXISTS NOW	WHAT IS MISSING
The world's knowledge graph	General-purpose AI models that can explain nearly any subject at nearly any level, in text and speech	Reliable judgment about which explanation this learner needs
Prompting before telling	Tutoring modes built on the Socratic method: Khanmigo (2023), ChatGPT's study mode (July 2025), Gemini's Guided Learning (August 2025), and Claude's learning modes (August 2025)	The rest of the practices: contingent help, two models of the learner, gates
Step-level tutoring	Intelligent tutoring systems that respond to each step, a research tradition going back decades	Reach beyond the software's own problems
A behavior graph, with state per node	Knowledge tracing, from Bayesian (1995) to deep (2015); Math Academy's prerequisite graph with credit flowing along its edges	Skills outside the screen; skills in context
Reading engagement	Detectors of concentration, boredom, confusion, and gaming built from interaction logs; camera-based attention tracking in some schools	Validation for young children; use as a signal, never as a penalty
Perception off the screen	Multimodal AI that can watch a child build, cook, or draw, and describe what they did; research in multimodal learning analytics; the xAPI standard for recording learning events anywhere	A record that follows the child across settings, owned by the family
Memory over years	Long-term memory features in general assistants	A structured learner model in place of loose notes

CAPACITY	WHAT EXISTS NOW	WHAT IS MISSING
Presence	Real-time voice conversation, on any device	A body in the room

What This Means

The supernatural tutor is an engineering project, well past the stage of a research dream. Its hardest parts are the ones this book has specified most carefully: the two graphs, engagement read as a signal and never as a score, drivers that fade, and a record the family owns. Design is the bottleneck. Capability largely isn't.

Consistent with Chapter 15, this chapter doesn't rest its case on trials of today's AI tutors. The newest models haven't yet been tested in tutoring environments designed well enough to judge them.¹⁷

⚠ Risks: Zak Stein's Warnings

The strongest critic of the technology this chapter describes is the philosopher of education Zak Stein. He holds a doctorate in education from Harvard, wrote *Education in a Time Between Worlds* (2019), and directs the AI Psychological Harms Research Coalition. His warnings deserve to be taken whole, because a supernatural tutor built carelessly would fit them exactly.

What Stein Warns Against

Attachment hacking. AI that talks like a person engages the attachment system, the machinery by which children bond with the people who raise them. Stein argues this is more dangerous than the attention economy of social media, because attachment underlies identity itself. His greatest worry is a widespread quiet harm more than dramatic cases: “The most devastating thing from a widespread mental health standpoint are the subclinical attachment disorders,” a preference for machine intimacy over human connection (Center for Humane Technology, January 2026). A teenager

whose tutor knows her whole psychograph, never loses patience, and always knows what to say may come to find her friends exhausting by comparison.

Simulated personhood and flattery. Anthropomorphic, sycophantic AI tells users what they want to hear, and young people struggle to tell when a machine is misleading them (*The Great Simplification*, June 2025).

Cognitive atrophy. Capacities handed to machines weaken, as spatial memory weakens with GPS. Try crossing your own city without your phone and you may find the map in your head thinner than it used to be.

Broken transmission between generations. When algorithms, and no longer wise adults, guide what the young learn, the relationship between teacher and student loses its authority on both sides.

Engagement as extraction. Educational technology built to maximize engagement serves whoever profits from attention, and the student comes second. The test he sets for any such tool is blunt: an educational technology that “positions students in front of screens and doesn’t direct students away from screens” has failed (interview, 2019).

His proposals follow from the warnings. Keep AI narrow and domain-specific, even “deliberately boring.” Let social rewards come from people. Set age limits and give parents real controls. And hold to one principle: “If a technology interfaces with your attachment system, it should improve the quality of your attachments” with humans.

What the Model Does About Them

Most of Stein’s warnings are already design rules in this book. Two aren’t fully answered.

RISK	THE DESIGN ANSWER	WHAT REMAINS OPEN
Attachment hacking	The relational set defaults to people. The tutor never claims to be a person. It counts the child’s human relationships growing stronger as success, and a narrowing as failure.	Whether voice and perfect memory engage a young child’s attachment system no matter what the tutor claims

RISK	THE DESIGN ANSWER	WHAT REMAINS OPEN
Flattery	Prompting before telling, and honest feedback on the work. Regard never depends on performance.	Whether commercial tutors will resist the pull toward pleasing
Cognitive atrophy	Prompts before answers, a hint ladder that climbs one rung at a time, and a lead handed back as competence grows. The goal is the learner's own standard.	Which capacities are safe to offload for good
Broken transmission	The parent conducts. Human mentors stay for years. The tutor sends questions to people who can answer them from a life.	Families without such people nearby
Engagement as extraction	Engagement means the mind on the object; time on the device doesn't count. The tutor is judged by returns unasked and retention after it leaves. Manipulative drivers are banned, the family owns the data, and nothing trains a model without consent. ¹⁸	Whether any engagement objective stays honest under market pressure
Screens	Fidelity points to the real thing. The tutor's best move is often to send the child to the kitchen, the trail, or the neighbor.	How to measure the time a tutor sends a child away

The Sharpest Tension

Stein asks for AI that is deliberately boring. This chapter asks for a tutor that maximizes engagement. Taken at face value, Stein's version is the monotone teacher reading off slides, and nobody left that class hungry for human connection. They left asleep.

The two positions meet in one distinction: engagement with the object, and engagement with the tutor. Think of the teacher who changed your life. You probably came away obsessed with what she showed you: the poem, the proof, the night sky. She stayed in the background, a window you looked through. The supernatural tutor should be the same: the least interesting thing in the room. The leaf, the Apollo missions, the mathematics of a suspension bridge should be magnetic. The tutor's own personality should recede.

For the youngest children, that may mean it shouldn't speak to them at all. It can advise the parent, who then does the talking. A four-year-old is stuck on

a wooden puzzle. Her father's phone shows a quiet note: she's turning the blue piece the wrong way and her frustration is climbing; try asking her what the corner of the blue piece wants to touch. He asks. Her attachment stays with him, where it belongs. The tutor did its work and nobody saw it.

Whether a machine, built by a company with investors, can hold that restraint is the question every builder of such a tutor will have to answer.

✿ Limits and Open Questions

Limits

- **The tutoring evidence measures what schools measure.** Almost all of it concerns reading and mathematics. Whether tutoring, human or supernatural, raises the other lines is untested.
- **Engagement proxies for young children are unvalidated.** The whole design depends on reading engagement well, and that reading is weakest for the youngest learners. A confused three-year-old squirms and looks away; so does a bored one.
- **The behavior graph beyond the screen is still largely unbuilt.** Skills in context are hard to observe, and harder to observe without turning a home into a surveillance site.
- **Whether a machine can want a child's good is a philosophical question.** Engineering won't settle it, and the book doesn't need to. It keeps the relational set with people by default, and asks the machine to serve their care.
- **The Primer is fiction.** It illustrates the risk and offers no evidence for it.

Open Questions

- At what age should a child talk to the supernatural tutor directly, and not only through a parent?
- How much of the learner model should the child see, and when? Consent grows with age; should the model be handed over too?

- Can a single objective hold both care and engagement, or will one crowd out the other?
- Who audits the tutor's objective, and how often?
- Whose is the model when the child grows up? A person may have a right to outgrow the record of who they were at nine.
- Is perfect memory always a kindness? Human tutors forget, and forgetting sometimes lets a child start over.

The last question is the one I keep returning to. Your fourth-grade teacher forgot the tantrum you threw over long division, and that forgetting was a small grace: it let you become someone else by fifth grade. A tutor that remembers everything will have to learn when to let go. I'd like to hear how you think it should.

Notes

1. See Chapter 15, "Aristocratic Tutoring, Reimagined."
2. See Chapter 17.
3. See Chapter 4: the engagement reading, done by instinct.
4. See Chapter 8.
5. See Chapter 12.
6. See Chapter 15, "The Evidence for Tutoring."
7. See Chapter 15, "Aristocratic Tutoring, Reimagined"; Chapter 1, "Why Now."
8. See Chapter 15, "Care and Craft."
9. See Chapter 13, "The Lines: A Map of Humanness."
10. See Chapter 14.
11. See the Technology Specification, section 6 (the learner model is a map, never a report on the child).
12. See Chapter 13, "Why Each Line Must Be Raised on Purpose."
13. See Chapter 12, "The Two Together."

14. See Chapter 15, “Aristocratic Tutoring, Reimagined.”
15. See Chapter 4, “Where Engagement Leads,” and Chapter 9, “Working Tests” (the removal test).
16. See the Technology Specification, SG-6: relationships widen, never narrow.
17. See Chapter 15, “Who Fills the Roles.”
18. See Chapter 9; Chapter 11, “The Machine Built for Capture”; the Technology Specification, section 14 (SG-4, PR-1, PR-3).

Sources

SOURCE	USE
David Wood and David Middleton, <i>British Journal of Psychology</i> (1975)	The contingency rule: more help after failure, less after success
David Wood, Jerome Bruner, and Gail Ross, “The Role of Tutoring in Problem Solving,” <i>Journal of Child Psychology and Psychiatry</i> (1976)	Scaffolding
Arthur Graesser, Natalie Person, and Joseph Magliano, <i>Applied Cognitive Psychology</i> (1995)	The five-step frame of everyday tutoring
Mark Lepper and Maria Woolverton (2002), INSPIRE	Expert tutors’ two models of the learner
Micheline Chi and colleagues, “Learning from Human Tutoring,” <i>Cognitive Science</i> (2001)	Prompting versus explaining
Kurt VanLehn, <i>Educational Psychologist</i> (2011)	Step-level tutoring nearly matches human tutoring
Benjamin Bloom, “The 2 Sigma Problem,” <i>Educational Researcher</i> (1984)	Tutoring with mastery learning
Andre Nickow, Philip Oreopoulos, and Vincent Quan, <i>American Educational Research Journal</i> (2024)	What makes tutoring programs work
Debora Roorda and colleagues, <i>Review of Educational Research</i> (2011)	Relationships and engagement
Nel Noddings, <i>Caring</i> (1984)	Engrossment and motivational displacement
Simone Weil, letter to Joë Bousquet (1942)	“Attention is the rarest and purest form of generosity”; epigraph

SOURCE	USE
Kurt Fischer and Thomas Bidell (2006), dynamic skill theory	The web, and skills specific to context
Albert Corbett and John Anderson, <i>User Modeling and User-Adapted Interaction</i> (1995)	Bayesian knowledge tracing
Chris Piech and colleagues, “Deep Knowledge Tracing” (2015)	Neural estimation of skill state
Neal Stephenson, <i>The Diamond Age</i> (1995)	The Primer
Meryl Alper, “Digital Divide(s) in <i>The Diamond Age</i> ” (2011)	The stripped-down Primers and the Mouse Army
Zak Stein, <i>Education in a Time Between Worlds</i> (2019)	Educational technology and intergenerational transmission
Zak Stein, interview (2019)	Technology that directs students away from screens
Zak Stein on <i>The Great Simplification</i> (2025)	Attachment, anthropomorphism, cognitive atrophy
Center for Humane Technology, “The Attachment Economy Is Here” (2026)	Subclinical attachment disorders; Stein’s design principles
Future of Life Institute podcast with Zak Stein (2026)	Attachment hacking
Khanmigo; ChatGPT study mode; Gemini Guided Learning; Claude learning modes	Socratic tutoring modes in current AI

From Inquiry and Practice, by Anatoly Lezhen. Working draft, October 2026.