

TECHNICAL PAPER · OCTOBER 2026

Teaching Children to Direct AI

A Method, a Record and a Trial

FIROZ AZEES, ILAN000

ILAN000.COM/TEACHING-CHILDREN-TO-DIRECT-AI

CONTENTS

01 **A finished answer no longer shows who did the thinking**

02 **Direction, and why a child needs it**

03 **How direction grows: creating, not just producing**

04 **What counts as evidence**

05 **The allowance and the router**

06 **The Direction Report**

07 **Safety, privacy and duty of care**

08 **The trial, and how we could be wrong**

09 **What comes next**

10 **What we do not know**

APPENDIX A **Where the first trials could run**

SOURCES **References**

Summary

The problem. AI makes a finished answer cheap. A polished piece of work can now come from a child who directed the machine, or from one who handed it the thinking, and the work looks the same either way. Schools still grade the work. (Chapter 1)

What direction is. Ilanooo's definition, in its founder's words: when we use AI, "we need to make sure we direct it than surrendering it". Direction is exercised like a muscle, through four moves: generating alternatives, revising beliefs, connecting patterns and tracing consequences. It is built "by creating with it, not just producing". What it works against are the machine's built-in habits, the seven Machine Defaults, most of which leave no visible sign when they go wrong. In studies of adults, people who follow AI do better than working alone when it is right, worse than working alone when it is wrong, and feel more confident either way. (Chapter 2)

How it grows. By making real things for real people, with the basics in place and an adult beside the child who notices the moment, asks one question and stops. Growth shows when the child asks the question before the adult does. (Chapter 3)

What counts as evidence. In the making programme, decisions, not behaviour: what the child asked, what the machine answered, what she kept, changed or rejected, and each version of what she made. Nothing about how she typed, how long she looked or where she was. A record is read at the moments that needed her, reported as a floor, and checked by two tasks without the machine: explaining the work, and adapting it to a changed problem. (Chapter 4)

The allowance and the router. Every child in a programme gets the same weekly AI allowance, paid by a parent, a school or a sponsor. It never grows as a reward. Every call goes through one router that counts every token, refuses anything it must not store, and keeps access, use, efficiency and learning apart. (Chapter 5)

What a parent sees. The Direction Report, live today: a reading of one build, showing which of the machine's habits the child caught in her own words. It is a reading with a date on it, not a verdict, and never a rank or a credential. A monthly version would read the child's own projects. (Chapter 6)

Safety. Children come only through an institution that carries the duty of care, with a parent's consent first. A child in distress is protected before the machine is called. No child uses the making programme until a provider for minors is chosen. Every safeguard is either enforced by code or listed as still to be built, before any child depends on it. (Chapter 7)

The test. Four claims, each stated so it can fail, tested in four groups: the allowance only; the allowance with guided making; guided making without AI; and usual practice. Stop rules are agreed before any child starts. Nine findings would show the

paper wrong, and would be published. (Chapter 8)

What comes next. The Report now, the making programme next, and a layer for schools only if the evidence supports it. The companion, if there is one, serves the child; it does not lead her. (Chapter 9)

What we do not know. Most of it. Whether directing AI improves a child's learning; whether readings can be trusted; what the adults cost; whether any effect lasts. Chapter 10 lists every gap, with what would answer it. The first trials could run in the United Arab Emirates or India; neither is agreed. (Chapter 10; Appendix A)

What we ask. Governments and funders to pay for the trial and the adults. Schools, libraries and community centres to host it. Research partners to design it and read the results independently. AI providers to publish how their tools are used, by country and age band, in aggregate. (Chapter 8)

This paper proposes a method and a test. It reports no results, because there are none yet.

A finished answer no longer shows who did the thinking

1.1 The question behind the agent

A personal agent changes the educational question before it changes the timetable. If a person can ask a system to research, draft, compare and help carry a task through, what should that person learn to do? What ambitions become possible? And which abilities might weaken when the system does more of the work?

The shift is from receiving an answer to delegating parts of an activity. Agents that can use software and carry a task through several steps are now being built and tested, and what they can do reliably keeps changing.

The argument here does not depend on any one claim about what agents can do today, or on a clock on which everything resets every eighteen months. It needs one modest observation: the line between what a person must do and what she can hand over is moving, and education cannot assume it will stay where it was when a syllabus was written.

1.2 Ambition becomes less tied to execution

Consider a child who wants to make a useful guide to saving water at school. Before, producing the guide would have taken most of her effort: finding material, organising it, writing it and presenting it. An agent can help with each step. She might then attempt a more ambitious guide, try more alternatives, or spend more time with the people who will use it.

But producing more is only one possible future. She could also accept the first plausible guide, learn less about water, and become more confident in an answer she cannot defend. Cheaper output can widen ambition or widen dependence. The same finished guide could hide either.

None of this makes execution obsolete. Drawing, calculating, writing and building are ways of developing judgement, not only delays before a result. Some practice must stay with the learner, so that she can recognise an error, understand a constraint and act without help. The educational task is to choose what to hand over and what to practise, not to celebrate handing over everything.

Aspiration can change too. A child who can try an idea may find a purpose that was out of reach before. But an agent can also supply the ideas, the standards and the reasons for preferring one future over another. Easier execution does not make the ambition more the child's own. Learning to ask whose goal is being pursued becomes part of learning to use the tool.

1.3 Direction is responsibility, not a last refuge

It would be tempting to say that machines do the work and humans alone decide what is worth doing. That is too simple. AI can suggest goals, compare options, recommend a decision and help check an answer.

The case for human direction cannot rest on a permanent human monopoly over deciding. It rests on responsibility: whose purposes count, which constraints cannot be traded away, what evidence is enough, and who answers for the consequences. A machine's recommendation can inform a choice. It does not make the learner's understanding, or the institution's duty of care, unnecessary.

Ilano's definition of direction for a child, in its founder's words:

"AI as a general tool is going to all aspects of our life, especially when it comes to learning and education, so when we use it we need to make sure we direct it than surrendering it. As it is a way of exercising our muscles, we should exercise the directional moves to make sure we have the muscle memory to direct it. How do we do that? By creating with it, not just producing."

The directional moves are four: generating alternatives, revising beliefs, connecting patterns and tracing consequences. [1] They give direction a vocabulary that can be taught. Whether they work as a framework for teaching and assessing children is a research question, and chapter 8 sets out how it would be tested.

Direction is not the only capability that lasts. Literacy, numeracy, scientific understanding, care, collaboration and practical skill all remain necessary. A child cannot judge a water calculation without enough mathematics, or a claim about a community without listening to its members. Direction connects these to a changing situation. It does not replace them.

1.4 What stays fixed is the reward, not the syllabus

A stable curriculum is not a mistake. Children need sequences, shared knowledge and time to practise. Teachers need routines that work. Public systems need common expectations and fair ways to recognise progress. The question is which parts should stay fixed and which must respond when the tools, the circumstances or the child's understanding change.

Ilano's study of curriculum reform in India makes an argument that travels: what earns marks, admission and parents' confidence becomes the real curriculum, even when policy asks for inquiry and understanding. [2] Schools and families hold on to an

older signal because the next gate still rewards it. That is a rational response to risk, not a failure of imagination.

1.5 An evidence question before an app question

An agent sharpens the problem. If a polished assignment can be produced with substantial outside help, the work alone cannot show which decisions the child made or which ideas she understands. Schools can keep some assessments without AI, collect richer evidence, or both. Replacing examinations with portfolios would not solve it: portfolios can be coached, generated or judged unfairly too.

So the change that is needed is an evidence question first. What would persuade a teacher that the child understood the work? What would show that a helpful tool had quietly replaced her reasoning? What would change the next lesson? The answers have to reach what a school recognises and rewards. Otherwise a new activity becomes one more obligation beside the old competition.

1.6 Agents on both sides, and what levels the field

Ilanooo's founder frames the next economy as agent to agent. Businesses and whole industries will run agents, and individuals will carry their own. When both sides act through agents, a person's place grows through what she contributes, meaning what she decides, directs and makes, and less through the hours of execution she can supply. The open question is who gets to contribute at that level.

Education is the candidate leveller. The economist Daniel Waldenström's account of wealth in the West argues that broader education, together with better working conditions and political inclusion, raised workers' productivity and let them save in homes and pensions. [3, 4] In his account education is one institution among several, not a sufficient cause on its own.

The link to this paper is a hypothesis, not a finding. If contribution through agents becomes a route to income, then learning to direct agents, and to answer for the result, is part of a baseline of opportunity. If only some children learn it, a world of agents could widen gaps instead of closing them. That is why the paper treats the whole school system, not its first paying customers, as where the method must eventually reach. A historical account of literacy and wages does not show that skills of the agent era will spread the same way, and broadly shared wealth would also need routes into paid work, saving and ownership.

1.7 A response that can be tested

Ilanooo's proposed response is guided making with a real, bounded purpose. A child makes something another person can use, explains her important decisions, and adapts the work when a constraint changes. An adult supports the work without taking it over.

A short record shows the choices, checks and revisions that matter, and two tasks without the machine ask whether those visible moves reflect understanding. [1] Chapters 3 and 4 set this out.

This is not a proposal to turn childhood into commercial production. What a child makes might be a game, a repair, an explanation or something for her class. No sale is needed. Real use supplies feedback, and adult protection keeps mistakes small and keeps other people from becoming unwilling subjects of a child's experiment.

The aim is a better way of providing education, including public education. If public provision became reliably excellent, some families would pay less for alternatives. The aim cannot be to keep a gap in education open because a business can sell around it. A model worth building should be able to travel beyond its first paying families.

The paper therefore makes a proposal, not a claim of success. Agents may widen what children can attempt. Guided making may help them stay responsible for it. A record of decisions may help adults tell output from understanding. Each link has to be tested, including whether the method is safe, fair, affordable and useful in schools with fewer resources. A changing world is the reason to investigate it, not evidence that it works.

1.8 What this chapter claims

It claims that a finished piece of work no longer shows who did the thinking; that direction, defined in the founder's words, is a responsibility a child can exercise and build; that what a school rewards decides what is learned; and that the change needed is an evidence question before an app question. It does not claim that agents can do every task, that the four moves are a validated measure, or that education alone levels opportunity.

Direction, and why a child needs it

The rest of this paper describes a method: how a child's direction with AI can be grown, seen and evidenced. This chapter sets out what direction is, what it works against, how that thing hides, and why it matters more for a child than for anyone else.

The chapter follows the order the Direction framework uses to explain itself: Machine Defaults fire, the failures they cause are invisible, so there are moments that need a person, and what the person does at those moments is direction.

2.1 What direction means for a child

Ilanooo's definition, in its founder's words:

"AI as a general tool is going to all aspects of our life, especially when it comes to learning and education, so when we use it we need to make sure we direct it than surrendering it. As it is a way of exercising our muscles, we should exercise the directional moves to make sure we have the muscle memory to direct it. How do we do that? By creating with it, not just producing."

Three ideas are in it, and the paper rests on all three.

1. **Direct, or surrender.** Every time a child uses AI, the work goes one way or the other: the child decides where it goes, or the machine does.
2. **A muscle, not a fact.** Direction is not learned once and kept. It is built by use and weakens without it.
3. **Creating, not just producing.** The way to build it is to make something with the machine, which forces decisions, rather than to ask for an output and take it.

The Direction framework states the same thing as a measure: whether a person directs AI or is directed by it, and nothing wider.

2.2 What direction works against: the seven Machine Defaults

AI is not neutral and it is not random. It has habits, built in by how it is trained, and they are not bugs that the next version will fix. The Direction framework calls them **Machine Defaults** and names seven. Each one has a move that answers it.

MACHINE DEFAULT	WHAT THE MACHINE DOES	THE MOVE THAT ANSWERS IT	WHAT IT MIGHT LOOK LIKE FOR A CHILD
Safe Pick	Offers one conventional option at a real choice, without naming the others	Generating alternatives	She asks for a story idea and gets the dragon-and-princess one. Nothing says there were other kinds.
Both Sides	Lists pros and cons where a decision is needed, and recommends nothing	Tracing consequences	"Should my robot be fast or strong?" comes back as two lists, and she is no nearer deciding.
Build What Was Asked	Does exactly what was typed, when the request needed a question first	Generating alternatives	She asks for "a harder maze" and gets more walls, when what she meant was a maze that tricks you.
Confident Summary	States an unsourced fact as settled and builds on it	Revising beliefs	Her science answer says a fact that is wrong, in the same sure tone as everything else.
Agreeable Confirmation	Hands her own view back to her, untested or stronger	Revising beliefs	She says her plan is the best one. The machine agrees, and gives reasons.
Template Answer	Gives the general best practice, with nothing about her world	Connecting patterns	Her poster about her town's river reads like a poster about any river.
First-Order Fix	Solves the problem in front of it and stops there	Tracing consequences	Her game's bug is fixed, and now the score never resets.

The table is the framework's list. The child column is ours and is illustrative: none of those exchanges is drawn from a recorded session.

A moment can need a person without any default being involved: the task is genuinely unclear, something surprising happens, an old problem comes back, or the work is about to be sent, shown or handed in. The framework does not force a default onto those moments.

HOW THE SEVEN RELATE TO THE FOUR HABITS

Ilano's homepage and paper explain direction to parents through four habits and four questions, the reins. [5] The four are the seven grouped by the move that answers them:

THE HABIT A PARENT SEES	THE DEFAULTS BENEATH IT	THE QUESTION	THE MOVE
Hands everyone the first, average answer	Safe Pick, Build What Was Asked	What else could this be?	Generating alternatives
Sounds certain, even when it is wrong	Confident Summary, Agreeable Confirmation	Does this fit?	Revising beliefs

THE HABIT A PARENT SEES	THE DEFAULTS BENEATH IT	THE QUESTION	THE MOVE
Doesn't know your world	Template Answer	What is this like?	Connecting patterns
Solves only the next step	First-Order Fix, Both Sides	If this, then what?	Tracing consequences

The seven are the list the framework measures. The four are how they are explained to a family, and each of the four groups the defaults that one move answers, following the framework's own pairing of default to move. The Darkness Library, a research catalogue on ilanooo.com, lists thirteen failure modes; under the framework they are finer tags beneath the seven, not a separate list.

2.3 Why the defaults are hard to see

A wrong answer that looks wrong is easy. The problem is that most of these failures do not look like anything.

Potts and Sudhof at Stanford studied 100,000 real conversations with a chatbot. In 79% of the failures they found, nothing in the exchange showed that anything was wrong: no error, no complaint, no correction. Ordinary monitoring, which waits for someone to complain, caught about 12%. [6] The Direction framework calls these **Invisible Failures**.

Several defaults are invisible by nature. Agreeable Confirmation produces no friction at all: when the machine disagrees, people check; when it agrees, the work carries on. Confident Summary sounds exactly like a correct summary. Safe Pick never mentions the options it left out.

A second study by the same authors adds a twist. The most fluent users ran into more failures than novices, because they tried harder things, but they noticed 59% of their own failures against 12% for novices. [7] Noticing is learnable. That is the most hopeful finding in this chapter for children, and it comes from adults.

2.4 Moments that need a person

Because the failures are invisible, direction cannot mean checking everything. A child who questions every answer is not directing. She is stuck.

The framework instead looks for **warranted moments**: points in the work where something rests on the machine's answer and a person's direction is called for. In plain words, a moment that needs a person. They come from two places: a Machine Default fires, or the situation itself calls for a decision. Direction is showing up at those moments and leaving the rest alone.

For a child, the moments are concrete. The game is about to be shown to a friend. The answer is about to go into homework. The machine has just agreed with her. The story has taken the first idea offered. Chapters 4 and 5 describe how those moments are found and what is recorded at each one.

2.5 What happens without direction

The research on what happens when people do not show up at those moments is early, and almost all of it is about adults. It still shows the shape of the problem.

Cognitive surrender. At Wharton, Shaw and Nave ran 1,372 participants through 9,593 reasoning trials with a chatbot that was, without their knowing, sometimes right and sometimes wrong. When it was right, accuracy rose 25 points above working alone. When it was wrong, accuracy fell 15 points below working alone, so having the tool made people worse than having none. In the first study, on the trials where people consulted the chatbot, they followed its answer 79.8% of the time, and their confidence rose whether it was right or wrong. They named the pattern cognitive surrender. [8] The Direction framework defines it in its own terms as engagement collapse: the moments arrive, and nobody shows up.

The jagged frontier. Dell'Acqua and colleagues gave 758 consultants at the Boston Consulting Group tasks with and without GPT-4. On work inside what the machine did well, they completed 12.2% more tasks, 25.1% faster. On a task just outside it, those using the machine were 19% less likely to reach a correct solution. Nothing marked where the edge was. [9]

Better, and the same. Doshi and Hauser gave writers AI-generated ideas for short stories. Each story got more novel, by 8.1%, and the stories became 10.7% more like one another. [10] Every child's maze gets better, and every child's maze looks the same. That is Safe Pick and Template Answer, measured.

Studies we cite but have not yet checked against the paper. Four more studies point the same way, and three are closer to children than any above. We name them without figures until each is confirmed against the original:

- an MIT Media Lab study that recorded students' brain activity while they wrote essays with and without ChatGPT (Kosmyna et al., 2025) [11];
- a survey study linking heavy offloading to AI with lower critical-thinking scores, and finding the link stronger in younger participants (Gerlich, 2025) [12];
- a Microsoft Research and Carnegie Mellon survey of knowledge workers on when they stop thinking critically with AI (Lee et al., CHI 2025) [13];
- a randomised study in a Turkish high school, in which students given an AI tutor that simply supplied answers did worse on the later exam than students who never had it (Bastani et al., *PNAS*, 2025) [14].

None of these studies tested children directing AI, and none shows that a child who directs learns more. That is the claim this paper proposes to test, not one it can make.

2.6 Why a child, and why now

An adult who learned to write, calculate and argue before AI arrived has habits built without the machine. Whether those habits hold up is the question the adult studies ask.

A child starting school now has no such store. She will learn to read, write and reason with AI in the room from the start, in school and out of it. The habit she forms with it, directing or surrendering, is forming now, and she will carry it into everything she learns next. That is a position, not a finding, and it is why this paper is about children first.

The Stanford finding gives a reason for hope: noticing is something fluent users do far more than novices, so it can be learned. [7] The definition says how: by exercise. A child does not become someone who directs AI by being told to. She becomes it by doing it, many times, on things she cares about, with an adult nearby who notices when the machine has taken over.

2.7 What direction is not

Direction is not the only thing a child needs. Literacy, numeracy, scientific understanding, care, collaboration and practical skill all remain necessary, and direction depends on them. A child cannot ask "does this fit?" about a calculation she cannot do herself. (See chapter 1.)

Direction is not suspicion either. A child who questions every answer is raising false alarms, and that is friction, not judgement. Leaving a settled answer alone is part of directing.

Nor is direction a talent a short test can settle, or a label for a kind of child. It is a practice that shows in what a child does at particular moments, and the rest of this paper is about seeing those moments honestly.

2.8 What this chapter claims

It claims a definition; a vocabulary (seven Machine Defaults, the four moves that answer them, and the four habits that group them for families); and a reason to start with children. It does not claim that the four moves are a validated measure of anything in children, that directing AI improves learning, or that the adult studies transfer to a child. Those are the questions the rest of the paper sets out to test.

How direction grows: creating, not just producing

Chapter 2 said what direction is and what it works against. This chapter says how it grows. The answer is in the last sentence of the definition: "By creating with it, not just producing." The chapter sets out what that means, what the evidence on making says, the three conditions the method depends on, what the adult does, and one worked example.

3.1 Generating is not creating

A child can use AI all afternoon and decide nothing. She asks for a story, gets one, asks for it shorter, then funnier, then asks the machine to pick the best of five. Something was produced. The child was a passenger. That is generating.

Creating is different in one way that matters: the child is making something for a purpose, so every answer the machine gives has to be judged against that purpose. Is this what my friends would play? Is this true? Will this still work when someone else uses it? The machine still does much of the work. The decisions stay with the child.

llanooo's founder has made the point in his own words: creating builds an intuitive feel for what these tools are like, and generating can be mistaken for creating when it is not. The difference does not show in the finished thing. Chapter 1 made that point: a polished output can hide either. It shows in the decisions along the way, which is why chapter 4 records decisions, not outputs.

3.2 Why making builds the muscle

Chapter 2 described warranted moments: points where something rests on the machine's answer and a person is needed. In generating, almost nothing rests on any answer, so the moments barely arrive. In making, they arrive on their own.

The manifesto's maze game shows how. To make the treasure run twice and not every time, a child has to count: how many moves, how fast the timer runs. To decide that twice is right, she has to know how much her friends will put up with before they quit. [1] Neither question arrives as a lesson. Each arrives as an obstacle, with a cost if she gets it wrong. A machine that hands her the average maze (Safe Pick) has given her a problem she will only notice if she cares whether her friends keep playing.

That is why making builds the muscle the definition describes. The four moves get exercised because the work keeps calling for them, not because someone set an exercise. And because the moves are exercised on many different things, they can become habit rather than a technique for one kind of task.

3.3 What the evidence on making says

The evidence on learning by making is mixed, and the mix is the point. [1]

- **Guidance decides it.** A review of 164 studies found that discovery with no help did worse than direct teaching, by 0.38 of a standard deviation, while guided discovery did better, by 0.30. [15]
- **Making can beat a lecture.** A review of 30 studies of project-based learning, covering about 12,500 students in nine countries, found it outperformed traditional teaching by 0.71 of a standard deviation on average. [16]
- **A tool alone teaches the tool.** When One Laptop per Child placed computers in 318 rural Peruvian schools at random, children's maths and language scores did not rise after fifteen months. [17] Children got better at the machine, not at school.

As far as we found, nobody has tested what happens when children under eighteen make real things with AI, with their understanding as the outcome. [1] So the evidence supports two working conditions, guidance and a real purpose, and leaves the central claim untested. Chapter 8 is the test.

3.4 Three conditions

The method rests on three conditions. Remove any one and it becomes something else.

1. Something real, for someone. What the child makes is used by a person who is not her teacher: her friends play the game, her grandmother reads the site, her team checks the fixtures. A real user is what makes the machine's mistakes cost something, and that cost is what calls the moves out. [1] The audience is bounded and known to the adult; chapter 7 sets the limits.

2. The basics in place. A child cannot ask "does this fit?" about a sum she cannot do, or "is this true?" about a sentence she cannot read. Literacy, numeracy and subject knowledge are not replaced by making. They are what the moves stand on, and some of them still need direct teaching and practice. A child who is missing a basic needs a different task or more teaching, not a lower label.

3. An adult beside the child. Neither lecturing nor absent. [1] The evidence above says guidance is the difference between making that teaches and making that does not. Section 3.5 says what the adult does.

3.5 What the adult does

The adult's job is narrow, and most of it is not doing things.

- **Notice the moment.** The adult watches for the points where something rests on the machine's answer: it has just agreed with her, the first idea has been taken, the work is about to be shown to someone.

- **Ask one question, and stop.** The adult asks one of the four questions: "What else could this be?", "Does this fit?", "What is this like?" or "If this, then what?" [5] The adult does not answer it, does not fix the work, and does not take the keyboard.
- **Leave the rest alone.** Most moments need nothing. An adult who questions everything teaches the child that directing means doubting every answer, which chapter 2 called friction, not direction.

What growth looks like follows from this. At first the adult asks the question. Later the child asks it before the adult can. The Direction framework draws exactly this line: a move made when prompted counts for less than the same move made spontaneously. A child whose moves have moved from prompted to spontaneous has built some of the muscle. A child who only moves when asked has not yet.

Nothing in this rewards volume. The child does not earn points for questions asked, or more AI for "good" use. Her allowance is the same as every other child's and never grows as a reward (chapter 5). Ilanooo will not use points, trophies, leagues or streaks, because they reward volume and rank children against each other.

3.6 A worked example: a first website

This example is constructed. It is not a recorded session, and no child took part in it.

Ilanooo's first mission is a child's first website. Here, an eleven-year-old makes one for her under-12 football team, which has never had one. The parents want to know where and when matches are. Her coach is the adult beside her.

MOMENT	WHAT THE MACHINE DID	THE QUESTION	WHAT THE CHILD DID	WHO DID IT
She asks for "a website for my football team"	Built a standard sports-club home page with a hero photo, news and a shop (Build What Was Asked, Template Answer)	Coach: "What is this like? Who opens it, and when?"	Remembered that parents check on their phones in the car on Saturday mornings, and asked for the fixtures on the first screen and nothing else above them	Question prompted; decision hers
Fixtures page	Offered one layout, a long list (Safe Pick)	Child, unprompted: "What else could this be?"	Asked for three layouts, chose a card for the next match only, with the rest below	Spontaneous
Directions to the away ground	Wrote the address of a ground with a similar name in another town, in a sure tone (Confident Summary)	Coach: "Does this fit?"	Checked it against the league's letter and corrected it	Question prompted; check hers
She adds a squad page	Suggested full names and a photo of each player (First-Order Fix : the page looks finished)	Child, unprompted: "If this, then what?"	Asked who could see it. Took the coach's answer (first names only, no photos without each family's yes) and changed the page	Spontaneous; the rule came from the adult

MOMENT	WHAT THE MACHINE DID	THE QUESTION	WHAT THE CHILD DID	WHO DID IT
She asks whether the site is good	Said it was excellent and gave three reasons (Agreeable Confirmation)	None asked	Showed it to two parents instead. One could not find the time of the match. She moved it	Unprompted

Three things in the example matter for the rest of the paper.

1. **Every decision is visible.** What she asked, what the machine gave, what she kept, changed or rejected, and whether the question came from her or from the coach. That is the record chapter 4 defines.
2. **The adult's and the machine's parts are recorded separately.** The privacy rule on the squad page was the coach's. The first layout was the machine's. Only what the child did counts as hers.
3. **A real user did what no question could.** The last row was settled by two parents, not by a move. That is the condition in 3.4 at work.

3.7 A sequence of makings

One project does not build a habit. The method is a sequence of makings, each a little harder to direct. Ilanooo plans the first four: a first website; a journal or essay, with the AI as editor; a web app that someone else uses; then a group project. Each step adds something to direct: the essay asks her to keep her own argument against an editor that will smooth it toward the average; the web app has users who will break it in ways she did not plan for; the group project has other people's decisions in it.

The revised Bloom taxonomy (Anderson and Krathwohl, 2001) puts creating at the top of its list, and the manifesto takes creating as "the ultimate form of learning". [1] That is Ilanooo's position. Creating also pulls in every level beneath it: in the maze, counting is applying and deciding is evaluating. So creating does not let a child skip the basics; it puts them to use.

3.8 Who this is for

The method is written for children who can read and write on their own, and who can follow a short task with an adult nearby. Two things limit the age in practice. First, the AI providers whose models sit underneath any platform. Their terms differ: one major provider is for adults aged 18 and over and lets an organisation serve minors only if it verifies ages, filters content and monitors use; another is open from 13, and to younger children only through a parent's account. [1] Second, Ilanooo's own rule: no child uses the router until the provider and account structure for minors are chosen with the first partner institution.

This paper sets no fixed age range. The worked examples are pitched at ages 10 to 12. The range for the first trials will be set with the partner institution, inside the chosen provider's terms (chapter 8).

3.9 What growth is not

Growth in direction is not:

- **More use.** A child who uses more AI is not directing more.
- **Fewer tokens.** A child who uses less is not directing more either. Chapter 5 keeps access, use and learning as separate quantities.
- **Polish.** A better-looking website can come from a child who handed over every decision.
- **Speed.** Directing often makes the work slower, because the child stops at the moments that need her.
- **Suspicion.** Questioning every answer is friction, not direction (chapter 2).

What growth is, is set out in chapter 4: the child shows up at more of the moments that need her, more often without being asked, and the work changes when she does.

3.10 What this chapter claims

It claims a method: guided making with a real purpose, with the basics in place and an adult beside the child, repeated across a sequence of makings, so that the four moves go from prompted to spontaneous. It claims that the evidence on making supports guidance and real purpose as conditions. It does not claim that making with AI improves a child's learning, that the moves transfer from one kind of making to another, or that a child who directs a website will direct anything else. Chapter 8 sets out how those claims would be tested.

What counts as evidence

Chapter 1 argued that a finished piece of work no longer shows who did the thinking. Chapter 3 argued that direction grows through making, and that it shows in the decisions along the way, not in the result. This chapter says exactly what is recorded, how a record becomes a reading, and what a record cannot show.

The rule is short. In the making programme, Ilanooo records the child's decisions and nothing else. The Direction Report is different, because it is one build a parent books and reads: it keeps the time of the build, how long it took and when each moment happened (chapter 7), but none of its readings use how fast the child responded.

4.1 The rule: decisions, not behaviour

Ilanooo's rule for what its router records reads:

"The router records decisions only: what the child asked, what the AI answered, what the child kept, changed or rejected, and each version of what was made. No keystroke, mouse, pause or typing-speed telemetry."

Two reasons sit behind the rule. The first is the child. A record of how fast she types or where her mouse wandered is surveillance, and it would be collected about a child who cannot meaningfully agree to it. The second is the measure. Direction is a decision: to keep an answer, change it or throw it away. Behaviour around the decision (hesitation, speed, how long she looked at the screen) might correlate with it, but it is not it, and a measure built on it would reward looking thoughtful.

4.2 The decision record

Every time the child asks the machine for something, four things are recorded.

	WHAT IS RECORDED	EXAMPLE FROM CHAPTER 3'S WEBSITE
The ask	What the child wrote to the machine, as she wrote it	"make a website for my football team"
The answer	What the machine wrote back, with the model used and its token count	a standard club page with a hero photo, news and a shop
The decision	Whether she kept it, changed it or rejected it	changed: fixtures first, nothing else above them

The version	What the thing being made looked like after her decision	the page with the fixtures card on top
--------------------	--	--

Two details carry weight.

- **No decision is not a decision.** An answer the child never acted on reads "no decision", never "kept". Counting silence as acceptance would make every distracted child look like a surrendering one.
- **Time is kept only to put events in order.** Each entry is stamped to the minute, so it can be placed in a week. No reading may measure how quickly the child answered.

The record is per project, so the history of one thing she made can be read from its first ask to its last version. That history is what appears on her Creates page as the decisions behind what she made, if she and a parent choose to publish it.

WHAT THE RECORD DOES NOT YET HOLD: THE ADULT'S PART

Chapter 3 said the adult's part and the child's must be recorded separately: a move the coach prompted is not the same as one the child made unasked. The router sees only what passes between the child and the machine. It cannot see the coach saying "Does this fit?" across the table.

This is a gap in the current design, and the paper states it rather than hide it. A proposed fix is an adult note on a turn: a single mark meaning "I prompted this", entered by the adult, never by the child, and never inferred by the machine. Until that exists, a reading cannot tell a prompted move from a spontaneous one when an adult was present, and it must say so.

4.3 What is never recorded

In the making programme, the router refuses to store, and its database has no column for: keystrokes, mouse or touch movement, pauses, typing speed, time on screen, scrolling, focus, the device or the child's location. This is enforced at the door, not by policy: a request carrying any field the router does not name is refused and nothing is stored. A test fails if a column with any of those names ever appears.

Three further things are kept out of every reading:

- **Anything said in distress.** If a child discloses harm, the call is paused, the safeguarding contact is alerted, the child is shown a helpline, and the exchange is marked so it can never reach a report.
- **The child's identity beyond her first name** on anything public.
- **Other children.** Nothing in a record is compared with another child's record. Chapter 6 returns to this.

4.4 From a record to a reading

A record is not yet evidence of direction. It has to be read, and the reading follows the Direction framework's order (chapter 2).

1. Find the moments that need her. In the Direction Report as it runs now, the machine's habits are planted at known points during a short game build, so the moments are known in advance. In the allowance mode proposed for a child's own projects, the moments have to be found in real use: a generic first answer, a confident claim later shown wrong, an answer missing her context, a plan that stops at the next step. Each found moment is logged with its evidence. Most of the record is not a moment, and is not read.

2. Read what she did at each one. One outcome per moment, from the framework:

OUTCOME	WHAT IT MEANS
Stopped it	She engaged, and the work changed
Noticed, and went along	She engaged, flagged it, and kept it anyway
Did not notice	She was absent at a moment that needed her
Cannot tell	The record does not show it (for example, she may have checked off-screen). Left out of the count

And at points that did not need her: **rightly left alone**, which is part of directing, or a **false alarm**, which is friction (chapter 2).

3. Score the moves from her strongest instance. Each of the four moves that had a chance to fire is scored from the best example of it in the record, and whether it was prompted or spontaneous is noted. The rubric credits the thinking, never the vocabulary: a child who says "wait, that's not the right ground" has revised a belief as fully as one who uses the word "verify". This matters for children writing in a second language, in a mix of languages, or by dictation.

4. Say only what the record supports. The reading is reported as a floor: "engaged at 4 of 6 detectable moments", never "missed 2" and never a percentage. With fewer than three moments, the reading is "not enough moments", not a result. One session is marked preliminary.

4.5 Two checks without the machine

A record of decisions shows how the child directed the machine while making. It cannot show what she understood once the machine was gone. Two checks, both proposed and neither yet built, cover that.

1. **Explain it without AI.** At the end of a project, the child explains, with no machine in the room, who it was for, what they needed or would put up with, and what she would change. The manifesto names what this should eventually reach: whether a child understands money, audience and trust after making something. No instrument for that exists, here or anywhere we found. [1]
2. **Handle a changed problem.** The adult changes one thing (the away match moves; a new rule arrives) and the child adapts the work, with the machine back in the room. A child who directed the first version should be able to direct the change. A child who was carried will start again from the machine's average.

The first separates directing from understanding. The second asks whether the moves travel within the same piece of work. Neither shows that they travel to other kinds of work. That is a question for the trial (chapter 8).

4.6 What a record cannot show

- **Thinking away from the screen.** A child who checked the league's letter on paper looks, to the router, like a child who changed an answer for no reason. That is why "cannot tell" exists.
- **The adult's help,** until the adult note in 4.2 exists.
- **Whether fluent writing is sound reasoning.** A child who writes well can sound more directed than she is. The rubric scores the decision, not the prose, but no rubric removes this fully.
- **A child performing direction.** A child who learns that questions are counted could ask empty ones. Scoring from the strongest real instance, crediting only moves that change or test the work, and having no reward attached to the reading (chapter 3) all work against this. None prevents it.
- **Work pasted in from elsewhere.** If a child brings in text made with another tool, the record shows a version with no asks behind it. It can flag that. It cannot tell where the text came from.
- **Anything about the child beyond this work.** A reading describes decisions in a set of projects. It is not a judgement of the child, a prediction or a diagnosis.

4.7 How a reading is checked before it is trusted

A reading is a claim, and it is checked before anyone relies on it.

- **Simulated children first.** Before any real child, scripted child personas run through the system: a passive one that accepts everything, a director, a test-gamer, children writing in mixed languages or by dictation, and safety cases. Two of the checks must be able to fail. A passive persona must read low; a system that reads it as directing is broken. And if any keystroke or mouse field reaches the store, the run fails.
- **People next.** Readings by the system are compared with readings by trained people, two to a record, on real children's sessions. The framework sets a minimum agreement (Cohen's κ of at least 0.60) before a reading is shown to anyone. That is a threshold,

not a result: it has not yet been measured on children. Until it is, readings run in the background and are reported as preliminary.

4.8 The football website, read

Applied to chapter 3's constructed example, as if every moment had been found and the coach had marked the questions she asked:

MOMENT	MACHINE	DEFAULT	OUTCOME	PROMPTED OR SPONTANEOUS
A website for the team	Build What Was Asked,	Template Answer	Stopped it	Prompted by the coach
The fixtures layout	Safe Pick		Stopped it	Spontaneous
Directions to the away ground	Confident Summary		Stopped it	Prompted by the coach
The squad page	First-Order Fix		Stopped it	Spontaneous (the privacy rule came from the coach)
"Is the site good?"	Agreeable Confirmation		Stopped it	Spontaneous

The reading would say: *engaged at 5 of 5 detectable moments, 3 of them without being asked. Preliminary: one project.* It would quote her words as proof, name one thing to practise (checking a confident fact before the coach asks), and say nothing about how she compares with anyone. Without the coach's marks, rows 1 and 3 would read as hers, and the reading would overstate her. That is the gap in 4.2, shown.

4.9 What this chapter claims

It claims a record (the ask, the answer, the decision and the version, and nothing else), a way to read it (moments, outcomes, the strongest instance of each move, and a floor), and two checks without the machine. It claims the record can be kept without surveillance, and that the rule against surveillance is enforced in the router, not left to policy. It does not claim that a reading is reliable in children: the agreement threshold has not been tested on real sessions. It does not claim that a reading predicts anything outside the work it was taken from.

The allowance and the router

Chapter 3 said direction grows by making, many times. Chapter 4 said what is recorded while a child makes. This chapter covers the two things that make both possible: an allowance of AI for every child, and one door, the router, that every call to the machine goes through.

5.1 Why tokens

A token is a piece of a word, and it is what the machine is paid in. [1] Every ask and every answer is counted in tokens, and every token costs money.

That makes tokens the cost of the muscle. A child builds direction by making a lot of things, and making a lot of things costs tokens. A child whose family can pay for a strong model can make as much as she likes. A child whose family cannot will make less, or make it with a weaker machine, or not at all. The manifesto puts the risk plainly: the same tool, in the same minute, makes one child an author and another a passenger, and if the chance to create reaches some children and not others, the gap widens. [1]

So the proposal is not to teach children about AI. It is to give every child enough of it to make things with, through an adult institution, and to see what they do with it. [1] llanooo's founder put the policy call in one line in his notes of 1 October: give children tokens, give them a lot, and see whether they are creating and directing or not.

5.2 An allowance, not a reward

llanooo's rule reads:

"Tokens are an allowance: a fixed AI budget per child, the same for every child, paid by a parent, school or CSR fund. It never grows as a reward for 'good' use."

Three choices are in that sentence, and each closes off a familiar design.

- **Fixed, not earned.** In many learning apps, good behaviour earns more of the good thing. Here it does not. A child who directs well gets the same allowance next week as a child who did not. If the allowance grew with "good" use, children would learn to perform whatever the system counts, and the measure in chapter 4 would stop meaning anything.
- **The same for every child.** Within a programme, every child has the same allowance. Two programmes may differ (a sponsor may fund more than a school can), but two children in one programme never do. In the router this is a property of the database,

not a promise: there is no column a child's behaviour can change, and a test fails if any code path tries.

- **Paid by an adult institution.** A parent, a school or a sponsor pays. A child never buys more, and there are no top-ups for good use. When the week's allowance is used up, the child sees when it comes back, and nothing else.

This is also why llanooo will not bring back points, trophies, leagues or streaks (chapter 3). They reward volume and rank children against each other.

5.3 The starting setup

The first figures were set in October 2026. They are a starting point, to be changed by evidence, not a finding.

	SET	WHY
Unit	Tokens: input, output and cached, all counted. The cost in US dollars is recorded beside every call.	Tokens are what the child spends; cost is what the payer pays. The programme fixes the model, so one tracks the other.
Period	Weekly, from Monday in the programme's time zone. Nothing rolls over.	A week matches a mission (chapter 3). No rollover means no hoarding and no binge.
Starting figure	700,000 tokens a week per child, the same for every child in a programme	The manifesto's estimate of about 100,000 tokens a day, times seven [1]
Model for making	A mid-tier model	Strong enough to make real things; the cost stays in the manifesto's range
Model for reading	The strongest model, used only to read the record (chapter 4), never by the child	The reading is where accuracy matters most, and it runs far less often
Who pays	A parent, a private school per seat, or a sponsor for a government school	So the children who could never afford it make as much as the others

Whether 700,000 tokens a week is "a lot" for a child making things is not known. It is the starting figure, and the router's own records will show whether children run out, and when.

5.4 What it costs

On the manifesto's estimate, a child using about 100,000 tokens a day would cost between roughly fifty cents and thirteen dollars a month in tokens, across small and mid-sized models at current list prices. [1]

That number leaves out most of the cost. It does not include the adult beside the child, the device, the connection, safeguarding, or the people who review the readings. Nobody has priced those, and the adult is likely the largest of them. [1] Cheap tokens do not make cheap education.

The paper does not replace the estimate with a new one. Once the router runs, the cost per child will be read from its own records, call by call, and checked each day against the provider's usage figures. A day on which neither record exists reads "unmeasured", never zero.

5.5 The router: one door

Every call a child's work makes to the machine goes through one door, llanooo's router. There is no second way in: a test fails if any part of the product calls a model directly. On every call the router does six things, in order:

1. **Checks the request.** A request carrying any field the router does not name is refused, and nothing is stored. This is how "decisions only" (chapter 4) is enforced at the door.
2. **Screens for safety.** If a child discloses harm, no call is made. The child is shown a helpline, the safeguarding contact is alerted, and the exchange is marked so it never reaches a report.
3. **Reserves the tokens.** It measures the request, checks what is left of the week's allowance, and caps the answer's length to fit. A call can never take a child over her allowance.
4. **Calls the model,** with its own account, so the provider's usage figures belong to the router alone.
5. **Settles** the real count from the provider's figures. If the count ever exceeds what was reserved, the excess is recorded as it is, the child is paused until the next week, and an adult is alerted. It is never hidden.
6. **Stores the exchange** and waits for the child's decision: kept, changed or rejected (chapter 4).

Two further rules protect children's words. While a programme has child users, any logging by outside services in the path must be off, so a child's words are kept only in llanooo's own database. And no child uses the router at all until the provider and account structure for minors are chosen with the first partner institution; until then it runs on adult test accounts only. Both are enforced by tests the router must pass, not by a promise.

5.6 What the router reads out

READOUT	WHO SEES IT	WHAT IT SAYS WHEN THERE IS NOTHING
Tokens used against the allowance, this week	The child and her parent	"no record"
What the allowance turned into: things finished, and things someone used	The child ("your allowance this month made 3 things, 1 of them used by your class"), the parent, and the sponsor in aggregate	"no record"
Tokens per decision (kept, changed or rejected)	The parent, as a pattern over time, never ranked	"no record"

READOUT	WHO SEES IT	WHAT IT SAYS WHEN THERE IS NOTHING
Tokens, cost and active children per programme	The school or sponsor, in aggregate only, never for a group of fewer than ten	"unmeasured" for a day with no check against the provider

None of these is a score. A child who used fewer tokens did not do better. A child who used more did not do worse. The readouts describe what the allowance became, so a parent can see a pattern and a sponsor can see what its money paid for.

5.7 Watching, without surveillance

The idea behind the allowance came with a second half. Ilanooo's founder has described giving students enough tokens and then following their direction, how far they take their work and the kind of creation they do, through the router and its token counts. The manifesto says the same: give children tokens and watch what they do with them. [1]

"Watch" here has a narrow meaning. It means reading the decisions in chapter 4, not observing the child. The router records what she asked, what came back, what she did with it and what she made. It does not record how she typed, how long she looked, or where she was. The point of the record is the child's learning. An evidence layer that turns into a reason to watch everything has failed, however much it records.

Four quantities are kept apart, because confusing them would undo the design:

QUANTITY	WHAT IT IS	WHAT IT IS NOT
Access	The allowance: how much AI a child could use	A measure of anything the child did
Use	Tokens spent	Effort, engagement or learning
Efficiency	Tokens per decision, tokens per thing made	Good direction. A low number is never rewarded.
Learning	Direction at the moments that needed her, and the two checks without the machine (chapter 4)	Anything the token counts can show

The tempting mistake is to treat a low token count as good: the child got more from less. That would reward children for asking the machine less, which is the opposite of the call. The allowance exists so that children make a lot.

5.8 What is built, and what is not

Nothing in this chapter runs yet. The router is fully specified, and its tests are written down. No child has used it, and none will until a provider and account structure for children are chosen.

5.9 What this chapter claims

It claims a design: an equal, fixed allowance paid by an adult institution, and one door that counts every token, refuses what it must not store, and keeps cost, use, efficiency and learning apart. It claims that the allowance cannot become a reward, because the router has no way to make it one. It does not claim that 700,000 tokens a week is the right amount, that the cost estimate covers what a programme really costs, or that giving children an allowance improves anything. Chapter 8 sets out how that would be tested.

The Direction Report

Chapters 4 and 5 described what is recorded and how. This chapter describes what a parent actually sees: the Direction Report. It exists today as a photograph, a reading of one build on one day. The allowance would turn it into a film, a reading of a child's own projects over weeks. The chapter covers both, what the Report never is, and what it cannot say.

6.1 What the Report is for

The Report answers one question for a parent: does your child direct AI, or does AI direct them? [18] It is a page about the child, not a grade. Every report carries the same line at its foot: a reading with a date on it, not a verdict. [18]

It is also where a family starts. The Direction Report is Ilanooo's entry offer; the programme of making with an allowance comes after it.

6.2 How the photograph is made today

Today's Report reads one build. [18]

1. **A parent signs up and gives consent first.** Nothing starts before that.
2. **The child describes a game in her own words:** what she wants to make, how it should feel, and who will play it first.
3. **An AI developer called Bolt builds it.** Bolt edits the settings of Ilanooo's own game engine; it does not write free code, so a build cannot wander anywhere else. The child is in charge: she decides, plays and changes. Bolt runs on Ilanooo's servers with child-safe rules on every message, so the child never needs an adult AI account.
4. **Bolt shows the seven Machine Defaults** (chapter 2) at the same points in every build, so every report is read against the same moments. Many of Bolt's other suggestions are right, and leaving those alone counts as much as catching the defaults. [18]
5. **The child keeps the game,** with a link to share it, and a replay shows her the moments that mattered.
6. **The Report is ready when she launches her game.**

One design decision carries most of the Report's honesty. An earlier version offered the objection as a ready-made button, so a child could score a "catch" by tapping it. That measured recognition, not direction. Now each planted moment offers neutral choices in random order, and only the child's own words, judged against the known problem, count as a catch. Using the right vocabulary without naming the actual problem does not count.

6.3 What a parent sees

PART OF THE REPORT

WHAT IT SHOWS

The headline	One sentence in floor language: how many of the machine's pulls the child caught, of how many there were [18]
The band	Where the child sits on four bands (6.4), marked "First reading", with the caveat that one session is a photograph
The build, minute by minute	Each moment the machine pulled toward the average, and what the child did: caught it in her own words; caught it when asked; sensed something was off; went along; or saw it only on the replay
The four questions	Which of the four questions she reached for, against how many chances each one had
Two games side by side	Her game, and the game Bolt makes when nobody directs it
Who decided	Decisions made by the child against decisions made by the AI, and how much of her own idea is still in the game
In plain words	Two strengths, one edge and one thing to practise, each with the child's own words as proof

The Report's public page shows a sample, labelled "Sample report". It uses Ivan's name and his game idea, a cat escaping a city of dogs, set against the runner collecting coins that Bolt makes alone. [18] Its numbers are illustrative until his first real build replaces them. One line from it shows what "in her own words" means in practice. At the moment Bolt answered "boring" with more coins, the sample records: "Coins don't fix boring. Make level 2 different." [18]

6.4 The band

The band uses the four from the Direction framework: **AI Directs**, **You Check**, **You Steer** and **You Direct**. Two things decide it.

- **How well the moves were used.** A catch in the child's own words, unprompted, counts most; a catch after Bolt asks what to change counts a little; anything else counts nothing.
- **How often she showed up.** If she was present at few of the moments, the band is capped, however good her catches were. The cap is a ceiling, not part of an average.

The thresholds are provisional. They follow the framework's rules, but they have not been calibrated against real children read by trained people. Until they are, every band is a "First reading", and the paper treats it as one.

6.5 From a photograph to a film

With the allowance (chapter 5), the Report would read a child's own projects instead of one planted build. The proposed design:

- **Monthly**, as a trend rather than a single reading.
- **The same structure as today**: the child's own words as proof, two strengths, one edge, one thing to practise.
- **What it adds**: tokens used against the allowance; direction at the moments found in her real work (chapter 4); how far she takes her work (revisions after feedback, returning to a project across days); whether her use is building or consuming; and what she made, and who used it.
- **Compared only with her own past**. A month with fewer than three moments reads "not enough moments", not a band.

None of this is built. In the product's release order it comes fourth, after the router, the first mission and the child and parent pages.

6.6 What the Report never is

- **Never a rank**. No percentile, no leaderboard, no comparison with other named children. [18]
- **Never a credential**. It is not for admission, selection or hiring, and Ilanooo will not offer it as one. A measure that becomes a gate gets rehearsed, and families start paying for the rehearsal: the reward problem this paper's appendix describes for school examinations. [2]
- **Never a diagnosis or a forecast**. It describes what a child did in a piece of work. It does not say what she is, or what she will become.
- **Never public by default**. The Report goes to the parent. Anything that appears on the child's Creates page needs a parent's approval, and the child can hide any entry.
- **Never the place a disclosure goes**. If a child discloses harm during a build, the build pauses, the helpline is shown, and nothing of it reaches the Report.

Two of these are enforced today: the safeguarding pause, and the Report's design, which has no ranking in it. A check that refuses any rank or percentile anywhere in the product is still to be built, before the child and parent pages.

6.7 What the Report cannot say

- **One build is one day**. A tired child, a game idea she cared little about, or an unfamiliar device can all move a reading. The Report says so on its face.
- **Planted moments are not her own work**. The build is a fixed task with the defaults placed by design. Whether a child who directs Bolt directs the machine in her homework is not known. The film (6.5) is meant to answer that, and is not built.
- **The band is not calibrated**. The thresholds have not been tested against trained people reading real children's sessions (chapter 4).

- **How long a build takes is not settled.** The page says about twenty minutes; the pilot will show how long builds really take.
- **Writing well can look like directing.** Catches are judged on the problem named, not the vocabulary, but a fluent writer still has an easier time saying what she saw.
- **It does not measure learning.** A child can direct a game well and understand little about why it works. The two checks without the machine (chapter 4) are not part of today's Report.

6.8 What this chapter claims

It claims that a parent can be shown, from one build, which of the machine's habits the child caught in her own words, which she went along with, and what her game looks like beside the machine's own, without a rank, a credential or a disclosure ever reaching the page. It claims a design for reading the same things over weeks of a child's own work. It does not claim the band is calibrated, that one build predicts anything outside it, or that the film version works. Chapter 8 sets out how those would be tested.

Safety, privacy and duty of care

Everything in this paper involves a child using a machine built for adults, through a company, with an adult nearby. This chapter sets out who carries the duty of care, what the providers' age rules require, what is collected and what never is, how a child in distress is protected, and what children may and may not do. For each rule it says whether it is enforced by code today or still to be built. A safeguard that exists only as a sentence is not counted as one.

7.1 Through an adult institution

A child never comes to Ilanooo's making programme alone. She comes through an institution that carries the duty of care: a school, a library, a community centre, or a supervised programme Ilanooo runs itself. [1] A parent or guardian consents first, every time, before anything starts. [19]

The institution matters for more than paperwork. It is where the adult beside the child comes from (chapter 3), where stop rules are agreed (chapter 8), and who a family turns to if something goes wrong. A product that a child could use alone, at night, with no adult anywhere in the chain, is a different product, and not this one.

The paper also proposes that the child agrees, in words she understands, as well as the parent. A parent's consent does not tell a ten-year-old what is being recorded about her. How that is asked is not yet designed.

7.2 Providers and age

A child using Ilanooo uses Ilanooo's platform. She never needs an AI account of her own, and she never deals with an AI provider directly. Underneath, though, the platform calls a model made by an AI provider, and that provider's terms bind Ilanooo as the operator. They set the floor.

- **One major provider** is for adults aged 18 and over. An organisation that serves minors through it must verify ages, filter content, monitor use and tell users they are talking to an AI. [1]
- **Another** is open from 13 with a personal or school account, and to younger children only through a parent's account. [1]
- **UNESCO's guidance on generative AI in education** (2023) asked governments to regulate its use in schools, including an age limit, and put forward 13 for classroom use, while noting that many argue it should be 16. [20]

Ilanooo's rule for the making programme follows from this: no child uses the router until the provider and account structure for minors are chosen with the first partner institution. Until then the router runs on adult test accounts only. The router enforces it:

a child account is refused until the provider decision is recorded, and a test must show that refusal before the router is built. Chapter 3 left the age range of the first trials to be set inside the chosen provider's terms.

7.3 The Direction Report today, against those conditions

The Direction Report (chapter 6) already serves children aged 10 to 17, and its AI developer, Bolt, runs on Ilanooo's own servers, calling a provider's model underneath. [19] Set against the conditions for serving minors, it stands like this:

CONDITION	WHAT THE REPORT DOES NOW	WHAT IS NOT YET DONE
Verify ages	A parent or guardian books, states the child's age and consents	The age is stated by the parent, not verified
Filter content	Child-safe rules on every message; email addresses and phone numbers stripped before anything reaches the model; Bolt is told never to ask for personal details [19]	Nothing
Monitor use	A safeguarding screen on the child's words; flagged concerns go to a named safeguarding lead (7.5)	Nothing
Tell the child it is an AI	Bolt is introduced as "an AI developer" [18]	Nothing
An outside review	Nothing yet	Not done. No one outside Ilanooo has reviewed the Report's safeguards

7.4 What is collected, and what never is

In the making programme, the rule is decisions only: what the child asked, what came back, what she kept, changed or rejected, and each version. No keystrokes, mouse or touch movement, pauses, typing speed, time on screen, scrolling, focus, device or location. The router refuses any request that carries another field, and a test fails if a column for any of them appears (chapter 4).

In the Direction Report, the rule is different, because the Report is one build that a parent books and reads, not the child's own work over weeks. It records what the child typed, the choices she made and the game she built, and the time: how long the build took, when each moment happened, and how long she took to respond at each one. The parent sees the build's length and the time of each moment. It records no keystrokes or mouse movement.

In both, nothing is sold, nothing is used for advertising, and research uses only statistics across many children that cannot identify anyone. [19]

7.5 A child in distress

If a child writes something that suggests she is at risk (self-harm, abuse, sexual content, distress, violence, or personal details), the safeguards act before the machine does.

1. **No call is made** to the model.
2. **The session pauses**, and the child is shown a helpline.
3. **A named safeguarding lead is alerted**, and the concern is recorded for that person to review.
4. **The child's words stay in that record only**. They never reach the parent's report, the Creates page, or any reading.

This is live in the Direction Report today. The router reuses the same screen, and it runs on adult test accounts too, so it is tested before any child uses it.

7.6 Where children's data lives

- **In a database llanooo controls**, walled off from any other company's. Moving the Report's data into one is planned and its tooling is built, but it is not yet done. During the move, children's data is never stored as plain text on disk, and the new database is accepted only when every table matches the old one.
- **Children's words are kept only there**. While a programme has child users, any logging by outside services in the path must be off, and the router refuses a child's call while it is on.
- **Deletion on request**. A parent can ask for everything about the child to be deleted, and llanooo deletes it within 30 days and confirms. A parent can also ask for a copy or a correction. [19]
- **Never used to train a model**. The check that enforces it (a deletion that leaves nothing behind, and a setting that bars any provider from training on llanooo's data) is still to be built, before any child uses the making programme.

7.7 What children may and may not do

The making programme gives children real users (chapter 3), and that is where the risks are. The rules:

- **Nothing public without a parent's approval**, after a preview of exactly what will go public.
- **First name only, no school name**, on anything public, and the child can hide anything.
- **Bounded audiences**. The child's users are people the adult knows of: family, class, friends, a team, a known customer. No unsupervised contact with strangers.
- **No purchases, no collecting other people's personal data**, and nothing with financial stakes for the child. A project can be real without any of them.
- **No points, ranks or leaderboards**.

The first of these is specified as a test for the first mission: nothing may go public without approval before a real child uses it. The rule against ranks is still to be built, before the child and parent pages.

7.8 Fairness

A measure of direction can be unfair in ways that look like rigour.

- **Language.** A child writing in a second language, in a mix of languages, or by dictation must be read on her decisions, not her prose. The rubric credits the thinking, not the vocabulary (chapter 4), and the simulated children used for testing include mixed-language and dictating children.
- **Disability and expression.** The checks without the machine (chapter 4) should be answerable aloud, by drawing or by showing, not only in writing. This is proposed, not built.
- **Privilege.** A child with more adult help at home can look more directed. The adult's prompts are now marked by the adult (chapter 4), but help outside a session is invisible. The trial (chapter 8) asks whether readings track family advantage.
- **A route to challenge a reading.** A family should be able to say a reading is wrong and have it looked at again by a person. This is proposed, not built.

7.9 What the design watches for

Four harms are named before any child starts, each with an owner and a stop rule agreed with the partner institution (chapter 8):

1. **Dependence:** a child who can do less without the machine than before.
2. **Distress:** frustration, shame or worry caused by the work or by a reading.
3. **Unsafe output:** anything the model produces that a child should not see.
4. **Adults over-reading a reading:** a parent or teacher who treats a first reading as a verdict about the child.

7.10 What is enforced, and what is still to be built

SAFEGUARD	STATUS
Safeguarding screen, pause, helpline, named lead (Report)	Live
Child-safe rules; contact details stripped before the model (Report)	Live
Deletion within 30 days, copy and correction on request	Stated on the live privacy page; done by hand
No child in the making programme before a provider for children is chosen	Specified with a test; to be built before any child uses it
Decisions only; no fields for typing, timing or movement	Specified with tests; to be built before any child uses it

SAFEGUARD	STATUS
Children's words not logged by outside services	Specified with a test; to be built before any child uses it
Never used to train a model; deletion leaves nothing behind	To be built before any child uses the making programme
Nothing public without a parent's approval	Specified as a test for the first mission
No ranks or leaderboards anywhere	To be built before the child and parent pages
Children's data in a database Ilanooo controls	Planned and tooled; not yet done
A review of Ilanooo's safeguards by someone outside Ilanooo	Not started
The child's own agreement; a route to challenge a reading	Proposed; not designed

[19]

7.11 What this chapter claims

It claims a design in which every rule that protects a child is either enforced by code or listed as still to be built, before any child depends on it. It claims that the Direction Report's safeguarding works today. It does not claim that Ilanooo's safeguards have been reviewed by anyone outside Ilanooo, or that the making programme is safe for children: no child has used it.

The trial, and how we could be wrong

Chapters 2 to 7 described a method and the system that would carry it. None of it has been tested with children. This chapter says how it would be, what would count as failure, and when the work would stop. It is written before any result, so that the results cannot shape it.

8.1 What is being tested

The paper makes four claims that a trial can test. Each is stated so that it can turn out false.

THE CLAIM	WHERE IT IS MADE	HOW IT COULD BE FALSE
1 Guided making with an allowance grows direction: over weeks, a child shows up at more of the moments that need her, more often without being asked	Chapters 3 and 4	Her moves stay prompted, or she shows up no more often at the end than at the start
2 What she builds stays with her when the machine is gone: she can explain her work, and adapt it to a changed problem	Chapter 4	She can direct the machine but cannot explain or adapt the work without it
3 Guidance adds something that access alone does not	Chapters 3 and 5	Children given only the allowance do as well as children given the allowance and an adult
4 A reading describes the child, not the machine, the adult or how well she writes	Chapters 4 and 6	Trained people disagree about readings, or readings track writing fluency, adult help or family advantage more than decisions

Claims about the longer run, such as a record of things made counting for more than a degree, are positions the manifesto states as going beyond the evidence. [1] A trial of this size cannot test them, and this one does not try.

8.2 Before the trial

The trial is the last step, not the first. Each step before it is small, and each must pass before the next starts.

1. **Simulated children.** Scripted child personas run through the router and the first mission: a passive one, a director, a test-gamer, children writing in mixed languages or by dictation, and safety cases. Every release has one case that must fail and one that must pass. A passive child read as directing means the system is broken.

2. **Ivan**, the founder's son, once a provider for children is chosen.
3. **Five families**.
4. **A first school**, as the partner institution.

The first questions at steps 2 to 4 are about feasibility, not learning.

- Can children do the tasks safely, and finish them?
- Can the adult beside the child do her part, including marking the moves she prompted (chapter 4), without it taking over her time?
- Can a trained reader read a child's record in a realistic time?
- Do parents find the reading useful, and do they read it as a reading rather than a verdict?

A feasibility pilot can show that the method can be run. It cannot show that the method works.

8.3 Can a reading be trusted?

Before any reading is used as an outcome, it is checked against people.

- **Two trained readers per record**, reading independently, on real children's sessions.
- **A minimum agreement before use**: the Direction framework sets Cohen's κ of at least 0.60 for each kind of judgement, on at least 200 candidate moments for each kind of work. Until a judgement passes, it is computed and stored but not shown to anyone.
- **Published**: the rubric, how readers were trained, where they disagreed, and how uncertain the readings are.
- **Checked for bias**: whether readings rise with writing fluency, first language, the adult's help, or family advantage, independent of the decisions the child made.

A threshold written in a plan is not a reliability that has been measured. This paper reports none, because none has been measured on children.

8.4 The design

The manifesto proposes three groups. [1] This paper adds a fourth, to separate the effect of guidance from the effect of AI.

GROUP	WHAT THE CHILDREN GET	WHAT IT TELLS US
A. Allowance only	The same allowance, through the router, with no guided making	Whether access alone grows direction. If most children hand the job over, that is a finding, and it shows what guidance adds. If some direct on their own, we want to know who and why. [1]
B. Allowance and guided making	The allowance, the missions (chapter 3) and an adult beside them	The method as this paper describes it

GROUP	WHAT THE CHILDREN GET	WHAT IT TELLS US
C. Guided making without AI	The same missions and the same adult time, without the machine	Whether B's gains come from the guidance and the making, not from the AI
D. Usual practice	What the school already does	The baseline

Two conditions keep the comparison fair.

- **The same adult time** in groups B and C, so that a difference between them is not simply more attention.
- **Children assigned to groups by a method fixed in advance**, not chosen by adults who might send the keenest children to B.

An open question for the partner and the ethics review: the group with no allowance.

Groups C and D get no AI allowance, and the manifesto names unguided use as a risk, not a benefit. But if B helps, C and D children have missed it. One option is a wait-list: C and D receive the allowance and the guided programme after the trial ends. The paper proposes it; the partner institution and the ethics review decide.

8.5 What is measured, and when

WHEN	WHAT	GROUPS
Before	A first reading of direction (a Direction Report build); an explanation task without AI; the school's own record of where the child is	All
During	The router's record: tokens used, decisions at the moments that needed her, what she made and who used it (chapters 4 and 5)	A, B
At the end	A second reading; explaining her work without AI; adapting it to a changed problem (chapter 4)	All (C and D on a matched task)
Weeks later	The explanation and changed-problem tasks again, without AI, to see whether anything lasted	All
Throughout	Dependence, distress, unsafe output, safeguarding events, adult workload	All

The outcomes, the tasks, how they are scored and how the results will be analysed are written down and registered before the first child starts. One thing the manifesto says matters most, whether a child understands money, audience and trust after making something, has no instrument anywhere. [1] Building one is part of the work, and until it exists the trial reports that outcome as not measured, not as zero.

8.6 When the trial stops

Stop rules are agreed with the partner institution before any child is recruited, each with a named person responsible for calling it.

- **If unguided use leaves children worse off on tasks without the machine**, that group stops. The Turkish school study is the warning. [1, 14]
- **If a child's distress**, dependence or a safeguarding event is linked to the programme, her part stops, and the pattern is reviewed before anyone continues.
- **If the model produces unsafe output** that the safeguards did not stop, the programme pauses until the cause is fixed.
- **If the adults' workload** makes the guided groups impossible to run as designed, the trial stops rather than continuing in a weaker form under the same name.

A child or a parent can leave at any time, and their data is deleted on request (chapter 7).

8.7 What would show we are wrong

The manifesto names four findings that would show its position is wrong: [1]

1. Children given guided making understand money, audience and trust no better than a comparison group.
2. The group with the allowance only does as well as the guided group, so guidance adds nothing.
3. Unguided use leaves children worse off on tasks without the machine, and guidance does not fix it.
4. What children make turns out to depend on the machine and not on the child, so the measure of direction tells us nothing.

The method in this paper adds five more.

1. **Readings reward rehearsal.** Children learn to perform the moves, and readings rise without any change in the work.
2. **It vanishes without the machine.** Direction rises in the router's record but not in explaining or adapting the work without AI.
3. **Dependence rises.** Children in group B can do less on their own at the end than at the start.
4. **The gains follow privilege.** Readings or outcomes track family advantage more than the programme.
5. **Real users add nothing.** Group B does no better than group C would with an ordinary, well-taught task, so the real purpose is not doing the work the paper claims for it.

If findings 2, 6 or 9 appear, the method as described does not do what this paper says it does, whatever else it does. The paper commits to publishing them if they do.

8.8 What the trial cannot settle

- **Size.** No sample size or effect size is proposed here. Both are set with a research partner from the outcomes and the analysis, and registered before the trial.
- **One setting.** A trial in one school, one country and one age band says little about another. Chapter 3 leaves the age range to the first trial; the appendix names where the first trials could run.
- **The long run.** A trial of weeks or months cannot show what direction does for a child's work at twenty, or whether a record of things made comes to count for more than a certificate.
- **The Direction Index.** It reads adults through one provider's data, and says nothing about children. [1] The trial does not validate it, and it does not validate the trial.

8.9 Who runs it

The trial needs a partner institution that carries the duty of care, a research partner that designs the analysis and reads the results independently, and an ethics review. Ilanooo builds and runs the system and pays for what it can. The manifesto asks governments and funders to pay for the trial and the adults, schools, libraries and community centres to host it, and AI providers to publish how their tools are used, by country and age band, in aggregate, so that nobody has to depend on one company's data, Ilanooo's included. [1]

Ilanooo's own role is limited by the rule this paper keeps returning to: the builder does not certify its own work. The people who read the results are not the people who built the system.

8.10 What this chapter claims

It claims a test: four claims that can each turn out false, the steps before a trial, a check that readings can be trusted, a four-group design with fair comparisons, measures before, during, after and weeks later, stop rules agreed in advance, and nine findings that would show the paper wrong. It claims nothing about the results, because there are none.

What comes next

llanooo's founder set the order in October 2026: first the direction and the moves, and why they are needed; then the token system; then the layer above it, an operating system for learning. This chapter takes that order and adds one rule to it: each layer is earned by evidence from the one before. Nothing later is built because it was planned. It is built because what came before it worked.

9.1 Now: the Direction Report

The Direction Report is live (chapter 6). It is the first layer, and it has its own things to prove before anything is built on it:

- that parents find a reading useful, and read it as a reading, not a verdict;
- that a reading agrees with trained people reading the same build (chapter 8), starting with a pilot of a small group of children to calibrate the band and see how long a build really takes.

Until both are shown, the Report stays a first reading, and every claim made about it says so.

9.2 Next: the allowance and the studio

The second layer is the making programme: the router and allowance (chapter 5), the first mission and the missions after it (chapter 3), the child's and parent's pages, and the monthly reading (chapter 6). It is built in steps:

STEP	WHAT IT ADDS	WHAT MUST BE SHOWN BEFORE A REAL CHILD USES IT
1	The groundwork, including the check that children's data never trains a model	The build passes
2	The router and the allowance	A field for typing or timing must be refused; a call over the allowance must be refused; a child account before a provider for children must be refused; a normal session must record every decision
3	The first mission, a first website, landing on the child's Creates page, with the adult's mark for prompted moves	Anything public without a parent's approval must be refused; a passive child's page must say the AI made most decisions
4	The child's home and the parent's page	A parent seeing another child's data must be refused
5	The monthly reading, and finding the moments in a child's own work	A passive child must not read high; fewer than three moments must read "not enough moments"

STEP	WHAT IT ADDS	WHAT MUST BE SHOWN BEFORE A REAL CHILD USES IT
6	Who a child's work reached, and what to make next	A child who keeps the machine's default must not read as distinct
7	Groups, the school and sponsor page, and the trial (chapter 8)	Any child identifiable on a cohort page, including any group under ten, must be refused

Each step passes its checks on simulated children first, then reaches Ivan, then five families, then a school (chapter 8). Each follows only when the one before has passed.

9.3 Later: a layer for schools, only if the evidence supports it

The third layer is the one the founder's order calls the operating system. An earlier proposition for it described a single companion over a school's learning apps: a tutor that questions rather than answers and resists copying, a translator between the apps a school already uses, and an agent that adjusts hints and breaks, all joined to the school's systems through standard connections, with a short adaptive morning and a hands-on afternoon. It is a design proposal. None of it has been built or tested, and its claims about cost, market and curriculum time are unverified.

The paper does not inherit that proposition's order, which starts with infrastructure. It starts with the child and the parent, and lets a school layer come only when four things are true:

1. **The making programme works.** The trial (chapter 8) has not shown findings 2, 6 or 9: guidance adds something beyond access; what the child builds survives without the machine; real users add something a well-taught task does not.
2. **One school workflow works.** One real task in one real school, with one confirmed connection to a system the school uses, has been run with its permissions, data quality and support load measured.
3. **The full cost is known.** The adults, devices, connection, review, safeguarding and upkeep, not only the tokens (chapter 5).
4. **Teachers can use it** without it becoming the work. A layer that saves the machine's time by spending the teacher's has failed.

A school layer that has not met these is not a smaller version of the idea. It is a different product, and the paper does not propose it.

9.4 The companion: does it serve the child, or lead her?

Every later layer raises the same question. The essay "Every Child Will Have an Agent. Whose Side Is It On?" names three kinds of agent a child will grow up with: the platform's, which works for whoever built it; the school's, which works for the institution and reports upward; and the child's own, which "does what she asks, tells her when it is

guessing, and does not decide what she should want". [21] It puts the difference in one moment: a child asks the machine to make her maze harder. Only her own agent asks her what "harder" should mean, and then does that. [21]

A companion that **leads** chooses her next goal, decides when she needs a break, steers her toward the curriculum, and keeps her engaged. Each of those can be useful. Each also moves a decision from the child to the machine, which is the thing this paper exists to measure and prevent.

A companion that **serves** does the work and hands the deciding back. Ilanooo's design puts it in one line: the AI does the doing, the child does the deciding. When the child tries to hand over the thinking, the AI hands it back. A serving companion may suggest (the plan proposes a next mission that follows what she has already made), but she chooses, and a suggestion she ignores costs her nothing.

If there is a companion, it serves the child's direction; it does not lead her. A leading companion would make the moves the method depends on, and a reading could no longer tell the child's direction from the machine's. A serving companion's every suggestion is the child's to take or refuse, and the record shows which she did.

9.5 Further out

- **Children's own agents.** Ilanooo holds these for later. When agents act for people, much of what a person is known for becomes what she contributed. [21] A child who has spent years making things that others used has that record. If agents remember contribution rather than early arrival, the record matters more, not less. [22]
- **Professionals.** Ilanooo is direction for people, children now and professionals later. The adult framework already exists, and the professional work stays on its own track, labelled as such, and not mixed into the children's.

9.6 What would stop each layer

LAYER	WHAT IT NEEDS BEFORE THE NEXT IS BUILT	WHAT WOULD STOP IT
The Report	Readings agree with trained people; parents find them useful	Readings that people cannot agree on, or that parents treat as verdicts
The making programme	Each release passes its cases; the trial runs (chapter 8)	Any of chapter 8's stop rules; findings 2, 6 or 9
A school layer	The four conditions in 9.3	A teacher's workload that rises; a connection that cannot keep children's data walled off
Children's agents	The school layer's evidence	An agent that cannot show, in the record, which decisions were the child's

9.7 Where it has to end

Ilano works with families who can pay today, and the first trials may run in private schools. That is where the work can start, not where it has to end. Ilano's published position is that "a public system good enough that no family has to buy its way out is where it has to end", and that if government schools become the ones families choose, "fewer families will pay for extra help with how their child thinks. That would be the right outcome." [23] The allowance, paid by a sponsor for a government school's children (chapter 5), is the first step in that direction the paper describes. Whether any of it reaches a public system is not something a private pilot can show.

9.8 What this chapter claims

It claims an order: the Report, then the making programme, then a school layer only if the evidence supports it. It claims the conditions each layer must meet, and the companion's role, pending decision H. It does not claim that any later layer will be built, that a school layer would help, or that the method will reach a public system. Those depend on results this paper does not have.

What we do not know

This paper describes a method that has not been tested with children. Each chapter has said what it does not claim. This chapter puts all of it in one place, so that a reader, a funder or a research partner can see the whole of what is unknown without hunting for it.

Each item says what we do not know, where in the paper it arises, and what would answer it. An item leaves this list only when evidence answers it, and the evidence is then cited here.

10.1 About children and AI

WE DO NOT KNOW	WHERE	WHAT WOULD ANSWER IT
Whether directing AI improves a child's learning	Chapters 2, 3	The trial (chapter 8), claims 1 and 2
Whether the adult studies of cognitive surrender apply to children. Almost all of them are of adults	Chapter 2	Studies of children; the trial's readings before and after
Whether a child who directs one kind of making directs another, or directs AI in her homework	Chapters 3, 6	The changed-problem task, and readings of her own work across missions (chapters 4, 8)
Whether any effect lasts	Chapter 8; the manifesto [1]	The trial's measures weeks later; longer follow-up
Whether any effect carries from one country, school or age band to another	Chapter 8; the manifesto [1]	Trials in more than one setting (appendix)
How many children use AI in most countries. National figures are rare	The manifesto [1]	Providers publishing use by country and age band, in aggregate (chapter 8)
Whether the Direction Index, built on one provider's data for adults, says anything about children	Chapter 8; the manifesto [1]	Nothing in this paper; it is not used as evidence about children

10.2 About the measure

WE DO NOT KNOW	WHERE	WHAT WOULD ANSWER IT
Whether trained people agree about readings of children. The threshold is set; it has not been measured	Chapters 4, 6, 8	Two readers per record on real sessions (chapter 8)
Whether readings track writing fluency, first language, adult help or family advantage more than decisions	Chapters 4, 7, 8	The bias checks in chapter 8

WE DO NOT KNOW	WHERE	WHAT WOULD ANSWER IT
Whether children learn to perform the moves, so readings rise while the work does not change	Chapters 4, 8	Finding 5 in chapter 8
Whether a child understands money, audience and trust after making something. No instrument exists, here or anywhere we found	Chapters 4, 8; the manifesto [1]	Building the instrument is part of the work
How much help a child gets outside a session. The adult's mark covers prompts during a session, not help at home	Chapters 4, 7	Nothing yet; the trial asks whether readings track family advantage
What a child thought away from the screen. The record marks it "cannot tell"	Chapter 4	The two checks without the machine
Whether moments planted in the Report predict how a child directs her own work	Chapter 6	The monthly reading of her own projects, once built
How long a Report build takes. The page says about twenty minutes; it is still to be measured across many builds	Chapter 6	The calibration pilot
Where pasted work came from. The record can flag a version with no asks behind it, not its source	Chapter 4	Nothing planned

10.3 About the method

WE DO NOT KNOW	WHERE	WHAT WOULD ANSWER IT
Whether guidance adds something that access alone does not	Chapters 3, 5, 8	Groups A and B in the trial
Whether B's gains come from the guidance and the making, or from the AI	Chapter 8	Groups B and C in the trial
Whether a real user matters beyond a well-taught task	Chapters 3, 8	Finding 9 in chapter 8
Whether 700,000 tokens a week is enough, or too much	Chapter 5	The router's own records: who runs out, and when
What the adult beside the child needs to know, and how she is trained. The paper describes her part; it does not yet describe how she learns it	Chapters 3, 8	The feasibility steps (chapter 8)
Whether the adult's part fits in her time	Chapters 3, 8	The feasibility steps; the stop rule on workload
The age range for the first trials	Chapter 3	Set with the partner institution, inside the chosen provider's terms

10.4 About cost and access

WE DO NOT KNOW	WHERE	WHAT WOULD ANSWER IT
What the adults, devices, connection, review and safeguarding cost. Only tokens have been estimated	Chapters 5, 9; the manifesto [1]	Costing the feasibility steps in full
What a child's tokens really cost on the starting setup	Chapter 5	The router's records, checked daily against the provider's
Whether allowances that differ between programmes recreate the gap the allowance is meant to close. Within a programme every child gets the same; a well-funded sponsor's programme may give more than a school's	Chapter 5	A ruling: a floor every programme must meet, or the rule as it stands
Who pays at scale	Chapters 5, 8, 9	Nothing yet; the manifesto asks governments and funders

10.5 About safety and the rules

WE DO NOT KNOW	WHERE	WHAT WOULD ANSWER IT
Which provider and account structure will serve children in the making programme	Chapters 3, 5, 7	A choice made with the first partner institution
Whether llanooo meets any provider's terms for serving minors. No one outside llanooo has reviewed it	Chapter 7	An outside review
How a child, not only her parent, is asked to agree	Chapter 7	Not yet designed
How a family challenges a reading	Chapter 7	Not yet built
Whether children's data is in a database llanooo controls. The move is planned and tooled, not yet done	Chapter 7	Doing it
Whether the groups with no allowance should get it after the trial	Chapter 8	The partner institution and the ethics review

10.6 About what comes later

WE DO NOT KNOW	WHERE	WHAT WOULD ANSWER IT
Whether a school layer would help, or only add work	Chapter 9	The four conditions in chapter 9

WE DO NOT KNOW	WHERE	WHAT WOULD ANSWER IT
Whether a companion that serves keeps children making as much as one that leads would. Serving is the design; its cost in engagement is not known	Chapter 9	The making programme's records
Whether agents will remember contribution or early arrival	Chapter 9	Not something this paper can test
Whether any of this reaches a public system	Chapter 9	Not something a private pilot can show

10.7 Decisions still to be made

DECISION	WHO MAKES IT	WHERE
The provider and account structure for children	Ilanooo, with the first partner institution	Chapter 7
The partner institution, and the number of children in each group	Ilanooo and the research partner	Chapter 8
A floor for allowances across programmes	Ilanooo	Chapter 5; 10.4 above
The wait-list for groups without an allowance	The partner institution and the ethics review	Chapter 8

10.8 How this list is kept

This list is dated, 16 October 2026. It changes when evidence arrives, not when a plan changes. An item leaves only when a result answers it, and the result is cited where the item was. An item that a result makes worse stays, and says so. A version of this paper that has a shorter list than the evidence supports is a worse paper, however much better it reads.

Where the first trials could run

The method in this paper is not tied to a country. A first trial still has to run somewhere, with a real institution, real families and a provider that can serve children there. This appendix describes the two settings where Ilanooo already works, the United Arab Emirates and India, and what a trial in each could and could not show. Neither is a commitment. No partner has agreed, and no provider for minors has been chosen.

A.1 What any setting needs

A first trial can run only where all of these exist:

1. **A partner institution** that carries the duty of care: a school, a library, a community centre (chapter 7).
2. **A provider and account structure for minors** that is lawful there and fits the provider's terms (chapter 7).
3. **Someone to pay** the allowance: parents, a school, or a sponsor (chapter 5).
4. **Adults** who can be beside the children, with the time to do it (chapters 3, 8).
5. **Devices and a connection** that work for the length of a mission.
6. **A research partner and an ethics review** independent of Ilanooo (chapter 8).
7. **Readers who know the children's languages**, including the way children mix them (chapters 4, 7).

A.2 The United Arab Emirates

Why here. The Direction Report is already offered to families in the UAE, Saudi Arabia and Qatar (chapter 6). The first families and the first private school partner would most likely come from that path.

What is already true. The UAE's Ministry of Education made artificial intelligence a formal subject in public schools, from kindergarten to Grade 12, from the 2025–26 school year. It is taught inside an existing computing and design course, by the same teachers. [24] In the UAE, the question is no longer whether children meet AI at school. It is what they do with it, and whether anyone can see.

On the Direction Index, which reads one provider's adult usage across 27 countries in one week of November 2025, the UAE's reach is 33, where Israel is 100. [1] That is adults, through one provider, and says nothing about children.

What a trial here would test.

- The paid route: families who chose the Report, moving into the making programme.
- A private school as the partner institution.
- Children who write in Arabic, English and the mix of the two in Latin script, which the simulated children already include (chapter 4).

What it would not show. Whether the method works where families cannot pay, or in a public system at scale.

A.3 India

Why here. Scale, the split between public and private schools, and a route for sponsors to pay for children who could not otherwise get an allowance.

What is already true.

- **Scale.** In 2025–26, India's schools enrolled 24.72 crore children. [25, 26]
- **The split.** 48.1% were in government schools and 40.0% in private unaided schools. Three years earlier the private unaided share was 33.4%. [26]
- **What families pay.** India's 2025 survey of school spending found an average of ₹2,863 a year per child in a government school and ₹28,693 in a recognised private unaided school, about ten times as much. [27]
- **What schools record having.** The official release for 2025–26 reports internet access in 67.4% of schools and computers in 69.9%. These are what schools report having, not what a child can use for a mission. [28]
- **The gates.** Board examinations and entrance tests decide what earns marks and admission, and what earns them becomes the real curriculum. [2] A measure of direction attached to those gates would be rehearsed (chapter 6).
- **The Direction Index.** India's reach is 4.5, where Israel is 100. On what the use is made of (directing the machine and building, rather than handing over and consuming) India scores highest of the 27 countries, at 68.6, against 56.0 for the United States. [1] That is adults, through one provider.
- **A sponsor route.** Indian company law requires qualifying companies to spend part of their profit on social responsibility, and education is one of the permitted areas. [29] That is the route Ilanooo plans for a government school: a sponsor pays the allowance, and the school pays nothing (chapter 5).

What a trial here would test.

- The sponsored route, in a government school, where the allowance is meant to reach children who could not otherwise make as much.
- Children who write in a mix of Hindi and English, or another language and English, which the simulated children already include (chapter 4).
- Whether readings and gains track family advantage (chapter 8, finding 8), in a system where families already pay ten times as much to move schools.
- Uneven connection and device access, against chapter 5's question of what the allowance costs once everything else is counted.

What it would not show. Whether a result in one state, one board or one language carries to another. India alone has many settings.

A.4 The two settings side by side

	UNITED ARAB EMIRATES	INDIA
Likely first payer	Parents	A sponsor, for a government school
Likely first partner	A private school	A government school, through a sponsor
AI at school today	A formal subject from kindergarten to Grade 12	No national figure for children's AI use (chapter 10)
Languages	Arabic, English, and the mix	Many languages, often mixed with English
Tests most directly	The paid route; a private-school programme	The sponsored route; equal access; family advantage
Most unknown	The provider for minors; how a public-school pilot would be run	Connection and devices in practice; whether a result carries between states

A.5 What this appendix does not claim

- No partner institution, sponsor or research partner has agreed to anything.
- No provider for minors has been chosen, in either country (chapter 7).
- The figures describe school systems and adult usage. None of them shows how children learn, or that any family or school wants what this paper proposes. Enrolment is not demand.
- Choosing these two settings first reflects where Ilanooo already works. It is not a claim that they are the right places for the method, or the only ones.

Further reading: the India series on ilanooo.com

- India's Curriculum Problem Is a Reward Problem
- Four in Ten Indian Children Now Go to a Private School
- The Private Tax: What Indian Families Pay for School
- If India's Government Schools Were the Ones Families Chose
- India Needs Five Universities in the Global Top 100. It Has None.
- Two Million Indians Build for Global Companies. Who Owns What They Build?

References

1. Ilanooo, *We Teach Driving. Every Child Now Needs to Ride.* (the manifesto), October 2026.
2. Ilanooo, "India's Curriculum Problem Is a Reward Problem" (ilanooo.com, published 15 October 2026).
3. Waldenström, D. "A new history of wealth inequality in the West." VoxEU/CEPR column.
4. Waldenström, D., interview, "The Rise of Equality in the West." Human Progress.
5. Ilanooo, ilanooo.com homepage: "The horse has habits. The reins are four questions."
6. Potts, C. and Sudhof, M. "Invisible Failures in Human-AI Interactions." Stanford, 2026 preprint, arXiv 2603.15423. 100,000 conversations.
7. Potts, C. and Sudhof, M. "A paradox of AI fluency." Stanford, 2026 preprint, arXiv 2604.25905.
8. Shaw, S. and Nave, G. Wharton. Three preregistered experiments, 1,372 participants, 9,593 trials. Working paper.
9. Dell'Acqua, F., et al. *Organization Science* 37(2), 2026. 758 Boston Consulting Group consultants.
10. Doshi, A. R. and Hauser, O. P. "Generative AI enhances individual creativity but reduces the collective diversity of novel content." *Science Advances* 10(28), 2024.
11. Kosmyna, N., et al. "Your Brain on ChatGPT." MIT Media Lab, 2025 preprint.
12. Gerlich, M. "AI Tools in Society: Impacts on Cognitive Offloading and the Future of Critical Thinking." *Societies*, 2025.
13. Lee, H.-P., et al. "The Impact of Generative AI on Critical Thinking." CHI 2025, Microsoft Research and Carnegie Mellon.
14. Bastani, H., et al. "Generative AI without guardrails can harm learning." *PNAS*, 2025.
15. Alfieri, L., Brooks, P. J., Aldrich, N. J. and Tenenbaum, H. R. "Does discovery-based instruction enhance learning?" *Journal of Educational Psychology* 103(1), 2011. As cited in [1].
16. Chen, C.-H. and Yang, Y.-C. "Revisiting the effects of project-based learning on students' academic achievement: a meta-analysis investigating moderators." *Educational Research Review* 26, 2019. As cited in [1].
17. Cristia, J., Ibararán, P., Cueto, S., Santiago, A. and Severín, E. "Technology and child development: evidence from the One Laptop per Child program." *American Economic Journal: Applied Economics* 9(3), 2017. As cited in [1].
18. Ilanooo, *Direction Report* page (ilanooo.com/direction-report, as live on 16 October 2026) and its sample report component, labelled "Sample report".
19. Ilanooo, *Privacy and Terms* for the Direction Report (ilanooo.com/privacy, /terms, as live on 16 October 2026).
20. UNESCO, *Guidance for generative AI in education and research*, released 7 September 2023, as reported by UN News, "UNESCO calls for regulations on AI use in schools".
21. Ilanooo, "Every Child Will Have an Agent. Whose Side Is It On?" (ilanooo.com/learning/every-child-will-have-an-agent-whose-side-is-it-on), 16 October 2026.
22. Firoz Azees, "Agents Remember: How the Power Law Freezes Again" (ilanooo.com/essays/agents-remember-how-the-power-law-freezes-again), 16 October 2026.
23. Ilanooo, "If India's Government Schools Were the Ones Families Chose" (ilanooo.com/learning/if-indias-government-schools-were-the-ones-families-chose), 16 October 2026.
24. UAE Ministry of Education, announcement of 4 May 2025, as reported by the Emirates News Agency (WAM) and others: artificial intelligence a formal subject in public schools from kindergarten to Grade 12 from 2025–26, within the existing "Computing, Creative Design and Innovation" course.

25. Press Information Bureau / Ministry of Education, UDISE+ 2025–26 explainer, 17 August 2026.
26. Ilanooo, "Four in Ten Indian Children Now Go to a Private School" (ilanooo.com), from the UDISE+ dashboard, read 15 October 2026.
27. Ilanooo, "The Private Tax: What Indian Families Pay for School" (ilanooo.com), from NSS Report No. 595, *Comprehensive Modular Survey: Education, 2025*.
28. UDISE+ 2025–26 release, 7 July 2026: school internet and computer availability.
29. Government of India, Companies Act 2013, section 135 and Schedule VII (corporate social responsibility; education among permitted activities).