

# The Framework

*A Plain-Language Guide to Aesthetic Accountability in AI Systems*

---

Something has changed about how we make things — and most people haven't noticed yet.

For most of human history, when you looked at something someone made — a painting, an essay, a building — you could learn things about the person who made it. Did they practice? Do they understand what they're talking about? Are they telling the truth about who they are? The thing they made was a window into them.

## **Generative AI has fogged that window.**

A 10-year-old can now type a prompt and receive a polished essay back. The essay looks impressive. But it tells you nothing about the child. A developer can generate an entire codebase without understanding what it does. A company can publish marketing copy that sounds human but carries no human judgment inside it. The output looks the same as it always did. The window is gone.

This is not just a technology problem. It is an accountability problem — one that shows up in classrooms, courtrooms, markets, and institutions every single day. And most of the people building AI systems were never given the tools to see it.

*This document is for them.*

## **The Broken Window**

Here is the simplest way to understand what has changed.

When a student writes an essay, the essay is evidence of their thinking. When an artist makes a painting, the painting is evidence of their choices, their practice, their perspective. When a programmer writes code, the code reflects their problem-solving judgment. We evaluate the thing to understand the person.

Generative AI breaks that relationship. Now a polished output can exist without the thinking behind it. The formal qualities remain — it looks finished, it sounds confident, it may even be technically correct — but it no longer reliably signals anything about the person who submitted

it. The connection between maker and made thing has been severed.

This matters in at least three places:

In education, assessment assumes that a student's output shows what they've learned. When AI generates that output instead, the assessment measures nothing about the student.

In law, copyright and authorship depend on tracing a creative act back to a responsible person. AI outputs have no such traceable source — the system isn't a legal person, and the user may not understand what the system produced.

In creative industries, we evaluate work based on who made it, what choices they made, and what their perspective contributes. When those things become indeterminate, the evaluation system breaks down.

**The question this framework asks is simple: if we can no longer judge the person by looking at the thing, what do we look at instead?**

## **The Mirror: Applying This to an AI System**

The clearest way to understand this framework is to apply it to the AI system you are probably already using: a large language model like the one that helped produce parts of this document.

Ask yourself four questions.

### **1. What did it learn from — and did those people agree?**

Large language models are trained on enormous amounts of human writing: books, articles, academic papers, websites, creative work. The people who wrote that material generally did not consent to having it used this way, were not compensated, and are not credited when the model generates something that draws on their style or ideas.

When the system produces something that sounds like a philosopher, or writes in the style of a novelist, it is drawing on work those people created — without their knowledge. That is a consent and labor problem built into the foundation of the system.

### **2. What does the user actually do — are they growing or skipping?**

This question matters most in educational and professional contexts. If a person uses the system to think more clearly — to push back on their ideas, sharpen their arguments, or understand something difficult — they are steering the process. Their judgment is still in the

work. The output reflects a collaboration where the human remained in charge.

But if a person uses the system to skip the hard part — to generate an essay they should have written, to produce code they don't understand, to create a strategy they haven't reasoned through — then the output belongs to no one. It evaluates nothing. It builds nothing in the person who submitted it.

*The same tool, used differently, produces completely different accountability outcomes.*

### **3. Who gets credit — and is that honest?**

When AI-generated content is submitted under a person's name, something about authorship becomes dishonest — not necessarily fraudulent in a legal sense, but misleading in a functional one. The person presenting the work is implying they made the judgments inside it. The system makes that claim impossible to verify.

A framework built around process-centered accountability asks: is there a way to make the production visible? Not to punish AI use, but to restore honesty about what happened and who is responsible for it.

### **4. What did it cost — and is that visible?**

Every AI interaction consumes real resources: electricity, water for cooling data centers, hardware manufactured under supply chains with their own labor and environmental costs. These costs are completely invisible in the interface. The experience is frictionless by design.

When material costs are hidden behind clean outputs, the people making decisions about AI use cannot weigh them. Accountability requires visibility. What is invisible cannot be evaluated.

## **The Shift: From Judging the Product to Evaluating the Process**

The traditional approach to evaluating anything — art, writing, software, research — focuses on the output. Is it good? Is it accurate? Is it sophisticated?

When AI is involved, that approach becomes unreliable. A bad thinker can produce a sophisticated output. A careless process can produce a technically accurate result. Judging the product no longer tells you what you need to know.

Process-centered accountability is the alternative. Instead of asking only "is this good?", it asks:

Was the process transparent about how this was made?

Was consent obtained from those whose work contributed to the system?

Did the person using the system develop their own judgment, or bypass it?

Are the material and environmental costs of the system proportional to its value?

These questions don't replace formal evaluation. Where it's clear that a human made the decisions — used AI as a tool the way a writer uses a dictionary, or an architect uses design software — traditional evaluation still works fine. Process accountability applies specifically where the human's role becomes ambiguous or absent.

## What This Framework Is Not Saying

To be precise about what this framework actually claims, it is worth being equally clear about what it does not claim.

**It does not say AI art is not real art.** It says that certain methods of evaluating art — ones that rely on reading the maker's choices through the work — become unreliable when AI is involved. The work may still have aesthetic properties. The evaluation method needs to adapt.

**It does not say using AI is wrong.** It says that how AI is used, by whom, under what conditions, and for what purposes — these things matter and need to be made visible. Legitimate, valuable, ethical AI use exists. The framework helps distinguish it from use that obscures accountability.

**It does not say human creativity is superior.** It says that where we use creativity as evidence of human judgment, skill, and perspective — in assessment, in authorship, in professional accountability — AI outputs do not provide that evidence in the same way. The comparison is not about quality. It is about what the output can be used to infer.

**It does not say these are solved problems.** The framework establishes the conceptual foundation. The specific standards, policies, and practices that should follow from it are open research questions — ones this work is designed to generate, not foreclose.

## Why It Matters Now

**There is a window closing.**

AI systems are being integrated into education, law, creative industries, and institutional governance right now — at a speed that outpaces the frameworks needed to evaluate them. The habits, expectations, and norms being established today will be difficult to revise once they

harden into practice.

If accountability is not built into these systems from the beginning — if the questions about consent, pedagogy, transparency, and material cost are not asked before deployment rather than after — several things risk being permanently foreclosed:

The capacity for institutions to evaluate whether people have developed actual judgment rather than access to impressive outputs.

The ability of creators — writers, artists, coders, researchers — to have their contributions traced, attributed, and compensated.

The possibility of designing systems where the process is as visible and evaluable as the output — before opacity becomes the default expectation.

This is not a warning against technology. It is an argument for building the right kind of technology — with accountability as a design condition, not an afterthought.

## How to Use This Document

This document is designed to be read by anyone who is building, deploying, evaluating, or regulating AI systems — regardless of their technical background or familiarity with aesthetics and philosophy.

It is not a checklist or a compliance standard. It is a set of questions to hold in mind when decisions are being made — questions that, if asked early enough, change what gets built and how.

The four questions introduced here — about consent in training, about whether users develop or bypass judgment, about honest authorship, and about material visibility — are diagnostic. They help you see the accountability structure of a system, not just its performance.

*What you do with that visibility is the work that follows.*