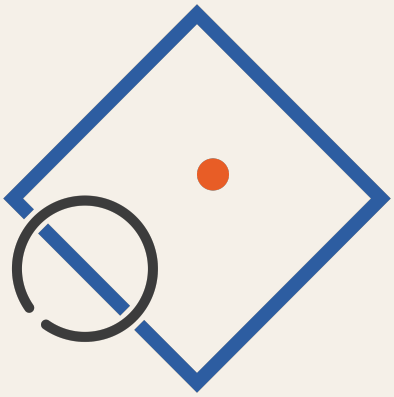

The Reference Is the Specification

revealed.design

Directing creative AI work by naming references instead of writing specifications

The Reference Is the Specification

Working paper, second edition – 2026

Summary

I built a portfolio site and its surrounding brand identity by directing an AI with cultural references rather than specifications. I did not write a brief. I named things – ILM, Panda Bear, a particular designer, a particular film studio – and the model unpacked each name into something implementable. This paper describes that method, why it works, where it came from, and where it fails.

The claim is narrow and I want to keep it narrow: when you direct an AI by naming a reference you know well, you are handing it a dense, pre-compressed specification. The work of creative direction moves from describing how to choosing what. As the cost of generating output falls toward zero, the binding constraint on creative production stops being budget or technical skill and becomes the quality of the references you can draw on – which is to say, taste. That is the whole argument. Everything else here is evidence for it or honest accounting of its limits.

This is a single-practitioner account. Sample size is one. It is offered as a usable method and a hypothesis, not as a proven result.

1. The claim

Conventional advice for prompting an AI assumes you should tell it how: specify the steps, the format, the constraints, the tone. The better your specification, the better the output. I found the opposite was the more useful move. When I stopped specifying and started naming – "make the motion feel like ILM," not "use a slow ease-out curve" – the output improved, and it improved in ways I could not have specified because I did not know they were possible.

A reference is a compression format. "ILM" is one word, but to anyone who knows the work it carries a motion philosophy, a standard of physical credibility, a particular relationship between weight and timing. A model trained on a large slice of the culture has that same unpacking available to it. When I say the word, it decompresses the word. I never wrote the easing curve. I named the studio whose easing curves I wanted to live near, and the model produced one.

This only works if you actually know the reference – not the name, the thing. You cannot name your way to taste you do not have. That is the catch, and it is also the point.

2. Why specification is hard

Photography is the cleanest precedent. When the camera arrived, it did not require the photographer to fabricate what the viewer would see. The camera recorded; the photographer decided where to point it, when to release the shutter, and which frame to keep. The craft was selection. The machine did the rest.

Photography did not kill painting. It killed painting's monopoly on representation, and painting went somewhere the camera could not follow. Every new instrument runs this pattern. The people who do well with it are not the ones who resist it and not the ones who surrender to it; they are the ones who understand that the tool changes what is possible without changing what matters. What matters is the judgment of what to keep.

Writing a full specification is hard for a reason that is easy to miss: most of what you want, you cannot articulate. You know the motion should feel a certain way. You do not know the easing function. You know the page should disclose itself in a particular order. You do not know the interaction pattern that produces that. A specification forces you to convert taste into instructions, and the conversion is lossy – you can only specify what you already know how to build. A reference does not require the conversion. It points at a finished example of the thing and lets the model do the decompression. The gap between the reference and the implementation is exactly where the model does its work, and keeping that gap open is what produces results you could not have ordered.

3. The reference as specification

The clearest evidence is the smallest. The core creative identity of the site came from roughly thirteen words of direction – proper nouns, not specifications. Each name resolved into something specific and implementable on the first attempt, without a clarifying round.

```
Reference · Domain · What it produced
ILM · Motion design · A specific easing curve
Pixar · Interaction timing · A settle-detection pattern (~300 ms)
Carrie Hobson · Content architecture · A three-perspective page structure
Riordan, animation · Data visualization · A Cobb–Douglas production-function graphic
Ingenue · Emotional register · Pacing and disclosure calibration
HAL 9000 · Collaborator identity · A named, persistent creative brief
Phoenix Standard Time · Time-aware behavior · A business-hours color shift
Panda Bear · Sonic identity · A layered chromatic synthesis
```

Figure 1. Eight references. The compression ratio – name in, full specification out – is the mechanism.

A representative full specification, for a page animation, ran to about thirty-five words: "at average reading rate, young ingenue 'enter' enters stage right from behind 'evil'; at first completely obscured, she remains on stage for the rest of the play." That is not an engineering spec. It is a theatrical direction. It names roles and a register, and it leaves the implementation entirely open. The animation that resulted is one I could not have written instructions for.

4. Three forcing functions

Naming things works because of three constraints the method imposes on the model. I call them forcing functions because each one removes a degree of freedom and forces the model to do work it would otherwise skip.

Compression. A reference is shorter than a specification, and the brevity is doing real work. A

paragraph-length brief invites the model to elaborate, hedge, and average. A single proper noun gives it nothing to pad – it has to resolve the name into specifics or fail visibly. Compression strips away the elaboration and forces the essential content.

Ambiguity. A reference under-determines the output. "ILM" does not say which easing curve. The model has to disambiguate, and disambiguation is where it adds information the brief never contained. A fully specified instruction forecloses this; an ambiguous reference requires it.

Directorial ignorance. I did not know what was technically possible. I did not know the relevant libraries existed or that a CSS transition could be staged like a theatrical entrance. That ignorance kept me from over-specifying. I could not collapse the work to a technical register because I did not have one. The direction stayed curatorial because curatorial direction was the only kind available to me.

These are not three good habits. They are one posture: name the destination and the feeling of arriving there, then leave the route to the instrument.

5. Constraints that emerge from the work

The three forcing functions all operate at the moment of direction. There is a fourth mechanism that runs the other way – backward, from the accumulated work onto the decisions that preceded it.

Partway through the project I rejected a material option – a set of gold-bordered sheets – without being able to say why. Asked later to articulate the basis, I produced a three-word rule: evidence over decoration. It had not been in any brief. It was not derived from a reference. It crystallized out of a single rejection.

What made it worth noticing is its timing. The rule arrived late, and yet the earlier work already conformed to it. The typography, the material choices, the refusal of decorative gradients – none of them had been evaluated against "evidence over decoration" when they were made, and all of them passed once it existed. The rule did not create the coherence. It named a coherence that was already there.

Emergent constraints have two properties worth recording. They are terse – "evidence over decoration" is three words – and the terseness is functional, not stylistic: a short rule survives being carried from one session to the next without being paraphrased into something softer, where a paragraph-length guideline drifts. And they are discovered, not designed. The honest version of this claim is modest: a constraint that explains earlier work without my having imposed it is suggestive evidence that the method has an internal logic. It is not proof. The proper test – does an independent reviewer find earlier work that violates the rule? – I have not run, and until someone does, "the earlier work all conforms" rests on my own assessment, which is exactly the kind of assessment an author should not be trusted to make alone.

6. What the method required

The method is portable; my results are not, and the gap between those two facts is the most important thing in this paper.

Three things had to be true at once for this project to happen. I had to be able to absorb the project's costs without noticing them – a subscription and about two thousand dollars in materials. I had to have a reference library worth drawing on, built over decades of paying attention to chairs, films, music, type. And the model

had to be good enough to unpack a proper noun on the first try. Any two of these without the third produces nothing: money and a capable model without taste produces polished work with nothing behind it; taste and a capable model without money never starts; taste and money without a capable model means hiring an agency.

Two of those three are getting cheaper. Subscriptions commoditize; models improve. The one that does not move is the reference library, because it only accumulates through lived experience and cannot be bought in bulk or prompted into being. So the method itself transfers – anyone can name a reference, and the transcript of how I did it is a usable template – but the output depends on the references behind it, and those are non-transferable. That is not a flaw in the method. It is the claim restated: if anyone could get these results, taste would not be the binding constraint, and the thesis would be false.

I should also be plain about the biography, because the original draft of this paper leaned on it heavily and a reader is right to be suspicious of that. Training in economics taught me to think in production functions and constraints; debate training taught me an economy of assertion; years of looking at design gave me the references. None of that is necessary to use the method. It is just the particular set of inputs I happened to bring, and I would rather state it once, plainly, than thread it through the argument as though it were proof of anything.

7. Limitations

This is one project by one person. Everything below is a real limit, not a rhetorical concession.

Sample size. n = 1. I cannot distinguish what generalizes from what is specific to me, this project, and one model at one moment. The method may depend on a particular kind of director, a particular domain, or a particular model's tendency to converge on aesthetic consistency. Other practitioners running the method is the only thing that would settle this, and that work does not exist yet.

The taste assumption. The entire method assumes my taste is real – that the references are well chosen and the judgments behind them are sound. If that assumption is false, the method does not fail halfway. It produces work that is internally consistent, technically correct, and aesthetically inert, and it produces it just as fluently. Nothing inside the project can test this. A fabrication vendor confirms that geometry is manufacturable; a deployment passes or fails; none of it speaks to whether the work is any good. That judgment can only come from an outside expert with domain knowledge who looks at the actual work. Until that has happened, the central premise of this paper is untested. This is the limitation I am least able to argue my way around, and I think it is the one a reader should weigh most heavily.

Attribution. Much of this paper's prose was drafted by the model under my direction, and the seam is not always visible. The model writes confident, well-formed sentences whether or not the underlying claim is sound – it cannot reliably tell its accurate output from its invented output, and neither can a reader going on style alone. That puts the burden of verification on me, and a reader has only my word that I carried it. Treat the factual claims here as claims, not as established findings.

Self-confirmation. Writing this the first time, I noticed a failure mode and want to name it against myself: a framework that reinterprets every observation as further evidence for itself is not strong, it is unfalsifiable. The honest test of this method is whether someone can run it and get a worse result than they would have

gotten otherwise, or look at this work and find it mediocre. I have tried to write this edition so those outcomes remain possible. If they are not possible, the paper is not making a claim.

8. Close

The established advice says the user must specify how. What I found is that naming what – precisely, from references you actually know – gets you further, and that the how emerges from the work. Knowing what you want turns out to be the scarcer and more valuable input.

That is a small claim. I have tried to keep it small on purpose, because the version of this paper that claimed more was the version that convinced no one but its author. The method is real and usable; whether the results are any good is a question I cannot answer from inside the project, and I would rather hand a reader that open question honestly than paper over it.

The reference is the specification. The taste behind the reference is the part that does not come in the box.

This is a working paper. The factual claims are the author's own account of a single project and should be read as such. Independent review – of the work, the method, and the claims – is invited and has not yet happened.

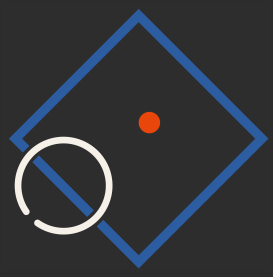

revealed.design

every stroke earns its place

made by Steven and SAL900X

May 2026