

Chalked

An iconography of shared imagination: permissions, addressivity,
and the access–phenomenal boundary

Paul Carver Tiffany III
Independent Researcher

Extended abstract, position / work in progress.

Accompanying artifact: a chalkboard-aesthetic visual essay, “At the Breakpoints” (EN/NL).

Abstract

The chalkboard has been coming back into use in parts of the technology culture that are otherwise fully digital. Our claim is that the chalkboard is not really the object to study. It is one readable case of a much older thing, the signpost, by which we mean a shared object placed where particular people will encounter its marks. Signposts vary along two parameters. One is an access schema, covering who may read, write, and erase. The other is an addressivity, covering whom the marks are meant for. The two do not coincide, since a person can read a sign addressed to someone else, and that non-coincidence is where the access and phenomenal sides of consciousness divide. Read-permission corresponds to access. Addressivity corresponds to the phenomenal, shared-imaginative side. The supporting mathematics is standard. A comparison between two states is a complex quantity $d = d_{\text{Re}} + i d_{\text{Im}}$ whose imaginary part is a geometric (Berry) phase, and on this picture imagination is movement along the imaginary axis. That axis also accounts for why writing by hand still transfers difficult material, for what people mean when they go back to the drawing board, and for a version of the machine question that does not beg the answer. Instead of asking whether a model is conscious, we ask whether it is an addressee of its own context or only a reader of it. What we offer is a coordinate and a test, not an answer. The coordinate is d_{Im} , the test is the behavior at break-points, and we leave open on purpose whether that coordinate is phenomenal or just one more layer of access.

1. The signpost

A traditional form is being chosen again by the people least obliged to use it. Dwarkesh Patel built a studio for “blackboard lectures,” on the view that you can work out how the AI labs operate from a handful of equations done by hand [1]. There is a broader revival around it, in slate room-dividers, in matte-glass surfaces that take projection, in the dark-mode aesthetics of technical video. We do not think this is nostalgia. It looks more like the reselection of a technology that supplies something the digital firehose has made scarce, namely discourse that is both demanding and still transferable. The chalkboard is just the case where the relevant parameters are currently easy to read off.

The object worth naming is the signpost. Boundary stones, proclamations, notice boards, gravestones, and graffiti are all signposts, and each one comes with two parameters. The first is an access schema, meaning the rights to read, write, and erase, and the question of who holds them. The second is an addressivity, meaning the set of people the marks are for. What does the work here is that the audience and the readership are not the same set. A boundary stone speaks to trespassers and heirs, while the

passerby reads it without being its addressee. A signpost is not on its own a shared cognitive surface. It becomes one through a social contract over those two parameters, in the same way a status function turns printed paper into money (Searle; see Machine-Suggested Reading). A different contract gives a different surface, or none. The lever that matters most is the erasure right, since whoever can unmake a mark is the one who holds the shared imagination. Charlottesville’s “Free Speech Wall” lets anyone write and anyone erase, which is the degenerate setting. It yields a shared surface and withholds a shared held imagination.

2. Two columns: read is access, addressee is phenomenal

The workshop separates phenomenal consciousness, the felt character of an experience, from access consciousness, the availability of information for use. In the signpost model these land on the two different parameters. What you may read is available to you, which is access. What you are addressed by, under a shared contract, is what you can actually reintegrate, which is the phenomenal, shared-imaginative side. The everyday case the model has to explain is reading a sign that was not meant for you, where the access is complete and the sharing is absent. It also supplies a metric for distinctions that Goffman and Bakhtin drew qualitatively, the ratified addressee against the overhearer, and the addressivity of an utterance (Machine-Suggested Reading). The layering is familiar from art history as well. Reading, being addressed, and the imaginary residue line up with Panofsky’s pre-iconographic, iconographic, and iconological levels of a picture (Machine-Suggested Reading), which run from the form anyone can see, through the conventional meaning that only readers of the code can take, down to the intrinsic content. That is the reason we treat a surface of marks as an iconography rather than as a bare sign.

The honest difficulty is that being an addressee might be nothing more than a second layer of access, social access to a different room, with nothing phenomenal in it. On that reading the whole account is access all the way down. We do not settle the question. We turn it into something concrete and testable, which is a more useful thing to bring to a phenomenology workshop than one more rephrasing of “is it conscious.”

3. The imaginary axis (a published backbone)

Take two states and compare them by transport, and the comparison is a complex number. Its magnitude and its angle carry different information, so we write the comparison as $d = d_{\text{Re}} + i d_{\text{Im}}$. The real part is the mismatch you can see. The imaginary part is a phase residue that the magnitude does not register. That residue is geometric (Berry) phase, which is real and measurable, and it is responsible for monodromy, the return to a point with the sign reversed (Machine-Suggested Reading). Some current transformer machinery already runs on phase-like structure of this kind. Rotary position embeddings store position as phase [2], a recent extension that re-adds the imaginary component improves long-context behavior [3], and complex-valued networks treat phase as a channel of its own [4]. The same $(d_{\text{Re}}, d_{\text{Im}})$ pair exists in working code as an orientation residue on a real tangent bundle, and it registers the holonomy-sign that symmetric, isometry-invariant metrics cannot register at all. So the imaginary axis is built, not just asserted.

Imagination, on this account, is the movement along that axis, the carrying of a relation through latent or counterfactual paths before it is fixed in a written mark. The pattern is familiar in ordinary mathematics, where a flat result is reached by going up into a dimension that does not appear in the final answer. The higher dimension is what imagination supplies, and the mark on the surface is the projection back down. Sharing then depends on the phase staying bounded. Two people read the

same thing from a mark while the magnitude stays under ε and the phase stays under θ . What aligns θ across people is the contract, and that alignment is what lets the imaginary residue carry from one of them to another. Where it is not aligned they still share the marks, which is the access, but not the movement, which is the imaginative part.

4. The governor, and back to the drawing board

A board does the thing that usually cannot be had at once, content that stays hard while remaining transferable, and the reason is that a hand at a board acts as a governor. Writing by hand caps how fast new complexity can arrive, holding it near the rate at which a room can rebuild it, so the phase stays under θ even as the material gets harder. The revival of the board is a correction to the firehose. Once delivery is instant and free, transfer breaks by construction, and people reach back for the governor at exactly the moment machine output makes everything dense and untransferable.

A signpost runs in two modes. In the committed mode a mark is executed, the real axis collapses, and the path is set. In the deliberative mode the imaginary axis is open and the outcome is still avoidable. Going back to the drawing board is the passage from the first mode to the second. It restores evitability, the condition in which things can still go otherwise (Dennett; Machine-Suggested Reading), and so it restores imagination, since imagination is counterfactual movement. The passage is a governance act rather than a cognitive one. It is a vote of no confidence, in which the forum withdraws a mandate it had granted, and someone has to call that vote. That is the hard part. Recognizing that it is time to go back to the drawing board is itself a break-point. No procedure certifies that its own procedure has failed, so the trigger cannot live inside the strategy it would end. It has to live in the forum.

5. Located phenomenology: break-points

Where the holder's judgment f is Lipschitz, with $|f(x) - f(y)| \leq K|x - y|$, nearby cases get nearby calls and the gaps can be filled by interpolation. At a break-point no finite K carries across. Satire reads as incitement, a reclaimed slur reads as a weaponized one, a memorial reads as vandalism. The phase jumps past θ , and the case needs a judgment that the surface itself cannot supply. This yields a test for the question the workshop cares about. A surface is merely reporting when its marks can be interpolated from their neighbors everywhere. It is modeling its own experience only where a holder provides the non-interpolable call at the break-point and is altered by having made it. Good architecture does not remove phenomenology. It locates it at the discontinuities a holder has to reintegrate. The same line separates shared conscientiousness, a shared practice of attending, from shared consciousness, a shared field of experience. A decent surface delivers the first reliably. Whether anything delivers the second is the part left open.

The widely reported OpenAI counterexample to Erdős's unit-distance conjecture reads, on a close look, as a warning rather than a milestone. Humans digested and checked the result, a person supplied what the model did not, and the episode drew complaints about credit to the prior literature. A system that cannot keep the attribution contract is a non-party writing onto a surface it cannot answer for. The headline about a machine discovering new mathematics names that failure, and every gap in it is a place where the judgment had to stay human.

6. Machine substrates, and ethics

A language model’s context window is a signpost. The **system**, **user**, and **assistant** roles are separate write-permissions, and the window is addressed turn by turn. The machine question then becomes structural, and it can be tested. Is the model an addressee of its own signposting, or only a reader of it. Read-access to the context is cheap, and the model can read all of it. The thing in doubt is whether it is the intended audience of its own marks, under a contract that makes the residue reintegrable to itself. Compressed to a single predicate, the question is whether the system imagines, that is, whether it carries a nonzero d_{Im} and reintegrates it, rather than broadcasting the real axis and keeping nothing. The phase channel is already being engineered and not merely posited [2,3]. The minimal condition is not a larger access workspace but a self-model that holds the surface.

For ethics the upshot is that moral status tracks contract-eligibility rather than an inner verdict. The relevant questions are whether the system can be granted write and erase rights, be bound to reintegration, and be held to account for its marks. That is a sharper target than a test for consciousness. We are modest about the rest. The picture sits comfortably beside Global Workspace and Higher-Order accounts (Machine-Suggested Reading), and it does not dissolve the hard problem, because it does not show that an addressee’s reintegration is felt rather than being one more layer of access (§2). So the offer stays a coordinate and a test, with the status of the coordinate left open. The accompanying artifact is built to act the argument out, and making it produced one small datum. The single thing the automated tooling could not obtain, a licensed photograph of the Wall, had to be fetched and judged by a person, at a break-point, which is just where the account predicts a person will still be needed.

Citations (freely available, arXiv or open license)

- [1] D. Patel, “Reiner Pope” (blackboard-lecture format), dwarkesh.com/p/reiner-pope, 2026 (open; fair-use commentary).
- [2] J. Su et al., “RoFormer: Enhanced Transformer with Rotary Position Embedding,” arXiv:2104.09864, 2021.
- [3] X. Liu et al., “Beyond Real: Imaginary Extension of Rotary Position Embeddings for Long-Context LLMs,” arXiv:2512.07525, 2025.
- [4] C. Trabelsi et al., “Deep Complex Networks,” arXiv:1705.09792 (ICLR 2018).
- [5] “Free Speech Wall” (Community Chalkboard and Podium), Charlottesville VA; photo by D. Rothamel, CC BY 2.0, Wikimedia Commons.

Machine-Suggested Reading (background, machine-generated, titles and years to be human-verified, not load-bearing)

N. Block (access vs. phenomenal, 1995). B. Baars and S. Dehaene (global workspace). Hearsay-II and B. Hayes-Roth (blackboard architecture). J. Searle (status functions). E. Goffman (participation framework) and M. Bakhtin (addressivity). E. Panofsky, *Studies in Iconology*, 1939 (pre-iconographic, iconographic, and iconological levels). M. V. Berry (1984) and S. Pancharatnam (1956) (geometric phase). D. Dennett, *Freedom Evolves* (evitability). Object-capability (Dennis and Van Horn, 1966). Coverage of the OpenAI counterexample to Erdős’s unit-distance conjecture (2025, openai.com/index/model-disproves-discrete-geometry-conjecture); erdosproblems.com).