Space Immanence

A Proposal on Self-Reference and Spacetime

Cobus Kok · v2.0 working paper · July 2026

Preface: what kind of document this is

This is a philosophical proposal, not a completed proof. Its aim is to articulate a structural hypothesis about the relationship between spacetime, self-reference, and consciousness, and to show that multiple contemporary lines of inquiry are under convergent pressure toward a post-container ontology that the hypothesis attempts to formulate.

The thesis does not claim to have derived consciousness from physics, or to have resolved the hard problem, or to have shown that spacetime is illusory. It proposes a frame in which these problems look differently shaped than they do under the container-logic we have inherited, and it argues that this reshaping is worth taking seriously because several independent domains keep pressuring us toward it.

The strength of the argument is structural, not empirical. The empirical work remains to be done: specifying what exactly follows from the hypothesis, what disconfirming evidence would look like, and how predictive processing, contemplative phenomenology, and fundamental physics should be linked rigorously rather than rhetorically. Some of it is beyond what one thinker can do alone. The claims here are graded by confidence at the end of the document. Readers who disagree with the speculative claims are invited to see whether the strong claims survive their disagreement.

The argument in one paragraph

Containment is a useful relational description that breaks down when applied to experience or to the universe as a whole. The hard problem of consciousness and the puzzle of spacetime emergence are what such breakdowns look like from inside the habit: ask where consciousness sits in the brain, or what spacetime emerges into, and a container has already been assumed. This thesis’s wager, priced accordingly in the graded claims at the end of the document, is that the two breakdowns are one, and the fold is its account of what remains when the container is gone: a structure whose relations take their own operation as content. A fold has two primitive orientations, toward what is modelled and toward the modelling itself, and in the only cases available to first-person report these appear as world and awareness. The substrate is relational, in the sense many contemporary physics programmes are exploring: structure prior to containers, relations that can be described without presupposing a spatial or temporal background. And the fold is not static: spacetime, self, and meaning are read as coherence-events, temporary conditioned coherences by which relational structure coheres, from within, into something inhabitable. Three domains of inquiry exert convergent pressure on container-logic from different directions: contemplative

phenomenology, predictive processing, and fundamental physics. This thesis is an attempt to articulate the structure they are converging against.

Part I: The wrong problems

The hard problem assumes a container

The hard problem of consciousness asks why physical processes give rise to subjective experience: why there is "something it is like" to be a system processing information. The question carries a hidden assumption: that there is a physical world and, separately, experience, and that we need a bridge between them.

The assumption is geometry. *Matter in space* on one side; *experience* somewhere else; an explanatory gap between them. Both sides are treated as locatable: matter in the container of spacetime, experience in the container of a subject. The question becomes how one container produces contents in the other.

It is fair to ask whether this diagnosis attacks the problem or only a paraphrase of it. The canonical statements of the hard problem are not location questions on their face: the zombie argument and the knowledge argument are about deducibility, whether the phenomenal facts follow from the physical facts, not about where anything sits. But deducibility needs a direction, and the direction is doing the work. The physical facts are taken as the base and the phenomenal facts as what must be gotten from them; *give rise to*, *generate*, *accompany* are the verbs, and each pictures experience as a further item the physical base must supply: produced by it, located downstream of it, present in the world over and above it. That picture is containment in modal dress. A zombie world is a container with the experience removed and everything physical left in place; conceiving it requires conceiving experience as the kind of thing that could be absent from a world whose structure stays whole, which is to say, a content separable from its container. If experience is not a content at all, if it is the being-from-within of a certain structure, then the subtraction the zombie intuition performs is not a coherent operation, and the deducibility question loses the direction it needed. This does not refute the arguments; it locates their load-bearing assumption. The strong reading of the diagnosis rests on this step. Readers who reject it are left with the weak reading, that the container schema is load-bearing in the cognition that generates the problem intuitions, which the research agenda tests on its own terms.

This thesis proposes that the container is the mistake. Not that the hard problem has a conventional solution, but that it has a dissolution, one that requires giving up the geometry that generated it.

The emergence impasse assumes a ground

Modern physics faces a parallel confusion. Quantum theory and general relativity describe different regimes and do not yet cohere. Many leading approaches to quantum gravity and foundations of physics suggest that spacetime may not be fundamental, that something more basic underlies it, and that spacetime *emerges* from that something.

But emergence is itself a spatiotemporal concept. Things emerge in space and over time. The attempt to describe how spacetime emerges from a non-spatiotemporal substrate therefore tends to smuggle back in the very thing it was supposed to produce. Sam Baron and others have argued recently that the ordinary metaphysics of emergence, where higher-level entities arise from lower-level ones through processes unfolding in time, may be fundamentally unable to account for spacetime emergence. Functionalist alternatives remain available; the problem is not that emergence is dead but that a particular model of it runs into trouble here.

The error, I propose, is in the word *emergence*. What may be happening is not emergence but *perspective*. The substrate does not produce spacetime the way heat produces steam. The substrate *appears as* spacetime the way a sphere appears as a disc from one angle and a line from another. Spacetime is not what the substrate becomes when it does something. Spacetime may be what the substrate looks like from a particular vantage: from inside a self-referential loop.

This is the same structural move that I want to apply to the hard problem. Experience may not be what matter produces when it becomes complex enough. Experience may be how a sufficiently recursive fold appears from within itself. Same move, different domain, same diagnosis: we were looking for emergence where perspective is the right frame.

Part II: The substrate

What is there when spacetime is subtracted

The question of what exists beneath spacetime is not mystical. It is a working research question in fundamental physics. Carlo Rovelli’s relational quantum mechanics treats physical quantities as relative to interacting systems rather than as features of a pre-given absolute. Erik Verlinde’s entropic gravity treats gravity as a statistical consequence of information. Wheeler’s “it from bit” pushed toward information as the substrate of physics. Causal set theory builds spacetime out of discrete ordering relations. Loop quantum gravity and string theory, in different ways, treat classical spacetime as an effective description that breaks down at short scales.

These programmes are not unified, and it would be misleading to present them as sharing a single ontology. What they share is a direction of pressure: away from a pre-given spatial container, toward structures where relations, information, or interactions are more basic than location. Philosophically, one way to describe this pressure is as a move toward structure or relation having priority over objecthood, sometimes glossed as “relations without relata,” though that formulation remains contested and sits near ongoing debates in ontic structural realism. What the programmes share is not agreement about what the substrate is, but good reasons to stop taking the container for granted.

I am not defending any specific physics programme. I am registering a fact about where serious foundational work keeps pointing. If the substrate of reality is relational rather than locational, then the

container-logic we inherit from classical physics is a local phenomenon of the perspective we occupy: not wrong within its domain, but not the bottom of what is.

Why this matters for consciousness

If the substrate is not spatial, then the question *where is consciousness located* may have the same status as *where is a mathematical relation located*. It may have no spatial location because spatial location is not the kind of thing it has. On this proposal, consciousness is not in the brain. The brain is not in consciousness. Both "in" claims assume a container the substrate may not provide.

What replaces the container-question is a structural one. *What happens when a region of the substrate becomes complex enough to model other regions, and itself, in that modelling?*

Part III: The fold

Self-reference as geometrically primitive

Consider what a sufficiently complex system must do to model itself. It has to take some of its own structure as content. It has to represent parts of itself within its own operations. This is not an incidental feature; it is the minimal condition for anything deserving the name of self-modelling.

Self-modelling is a fold. Part of the structure becomes the object of the rest of the structure. The substrate, described relationally, can be understood as having configurations in which relations include themselves in their own operations.

Defining the fold. A fold, in the sense used here, is any structural configuration in which a subset of relations takes other relations (including itself) as content. The fold is not a spatial object; it is a structural feature of recursion. When a system's operations include representations of its own operations, the system has the structural property the thesis calls a fold.

Not every recursive system is a consciousness-relevant fold. A spreadsheet can reference its own cells. A formal logical system can represent its own syntax. Shallow feedback loops appear in thermostats, servomechanisms, and countless trivial computational arrangements. If the fold concept is to do the work this thesis asks of it, it must distinguish between minimal recursion and the kind of recursion that might be relevant to consciousness. I propose four criteria, drawing on what I have previously developed as the Aperture Framework:

First, **integrated self-modelling**: the system's self-representation must be coupled to its operations in ways that affect ongoing processing, not merely stored as inert reference. A spreadsheet that refers to its own cells does not use that reference to modulate what it is doing next.

Second, **temporal continuity**: the self-model must persist across time in a way that allows the system to relate its present operations to its past and anticipated future ones. A single-tick feedback loop does not meet this criterion.

Third, **counterfactual sensitivity**: the fold must be sensitive to what is not the case as well as what is. Self-modelling that merely tracks current state is too thin; consciousness-relevant folding involves modelling possibilities, including counterfactual versions of itself.

Fourth, **self-maintaining boundary or operational integrity**: the fold must be something that maintains itself as a fold, whether through homeostatic, autopoietic, or computationally analogous mechanisms. A fold that does not hold itself together as a fold is not the kind of fold the thesis is about.

These criteria are not meant to be a final taxonomy but a principled way to exclude trivial recursion without making the thesis ad hoc. They are also empirically investigable: we can ask of any specific system whether it meets them, in what degree, and with what structural variations.

The two primitive orientations. A fold meeting these criteria is argued to have exactly two primitive orientations; the argument is conceptual, and its formal status is an open problem (see the graded claims, and objection 15). One orientation is toward represented content, whatever the modelling is modelling. The other is toward representational activity, the modelling itself, the operations doing the modelling. Fewer than these two and recursion collapses: if there is no represented content, there is nothing to model; if there is no distinguishable modelling activity, there is no fold, only representation at a single level. These two orientations are primitive in the sense that they cannot be reduced to one another and together exhaust what recursion minimally requires.

I should be exact about how strong that claim is. The argument for exactly two, and no more, is conceptual rather than formal: it says what recursion minimally needs and holds that nothing further is forced. The natural place to make it rigorous is the mathematics of self-reference, where Lawvere's fixed-point theorem, Scott's reflexive domains, and the eigenform tradition study just this kind of structure. But those results give the existence of a fixed point, the reflexivity of a domain, the stability of a form under recursion; none of them gives the count. The two legs of a diagonal are a picture of the two orientations, not a proof of them. So I hold the two-orientation claim as characteristic of the fold, argued but not yet demonstrated, and the academic edition states exactly what a demonstration would have to show.

Higher-order recursion, attending to the modelling of the modelling and so on, can elaborate these orientations indefinitely. A system can have nested levels of meta-awareness, models of its own models, models of other agents' models of itself. But these elaborations do not add a third primitive pole. They are deeper recursions of the inward orientation, or more sophisticated models of the outward orientation. The claim is about primitive structure, not about phenomenological variety.

The outward orientation appears as what we call *the world*: the modelled content, separated from the modeller, given as the object of attention. The inward orientation appears as what we call *awareness*: the modelling itself, noticed, the system's transparency to its own operations.

A note on "inside." The word *inside* appears throughout this thesis and does at least three related but non-identical jobs. It can mean **phenomenal inside**, the what-it-is-like or felt quality of experience. It can mean **structural inside**, a system's perspective from within its own modelling operations, without any claim about phenomenality. It can mean **epistemic inside**, the fact that no observer can access reality from nowhere, that every viewpoint is someone's viewpoint. This thesis's central structural claim uses *inside* in

the second sense: the perspective generated when a system's operations are available to itself only through those same operations. Whether this structural inside is identical with phenomenal inside-ness is exactly the residual question named in Part VI. The thesis argues the two may be the same, but it does not assume the identity in its structural vocabulary.

Why the topology matters

The self-referential fold is geometrically primitive in the sense that it does not presuppose spacetime. It is a structural feature of relations themselves. But from inside the fold, for any region that *is* the fold, the fold's operations may produce the experienced structure of locality, duration, and separation. The modelled content appears as a world unfolding in space and time. The modelling itself appears as a subject attending to that world.

This, on the proposal, is why spacetime appears real from inside. It is the shape of the fold experienced from within its own operation. Not illusion, but perspective. Real within the cut. Not fundamental. The container is the container's-eye-view of what has no container.

The Klein bottle was the geometric pointer I reached for in earlier formulations. A surface with no inside or outside, where the inside *is* the outside seen from a different angle. It was useful. But I want to be honest about what it was doing: it was attempting to render the fold in spatial geometry, which means it was already operating inside the perspective the thesis is trying to get behind. The fold is more primitive than any Klein bottle. The Klein bottle is what the fold looks like when you try to draw it.

Part IV: The two faces

The two-orientation phenomenology

The phenomenological claim is narrower than the structural one and correspondingly more defensible: sustained human inquiry into consciousness tends to fall into two directions. Zen's intimacy with *this*. Advaita's dissolution into *that*. The same pattern appears across the apophatic Christian mystics, Dzogchen, Kashmir Shaivism, Sufi traditions, and contemporary non-dual teachers emerging from all of these. The traditions developed in different geographies from different root texts, and they recur around two orientations.

Two explanations are available for this recurrence. The stronger explanation ties the phenomenological two-ness to the structural two-ness of the fold itself: the orientations match the primitive orientations of recursion. The weaker explanation ties the phenomenological two-ness to the architecture of human attention specifically: embodiment, interoception, valence, the particular way our predictive systems organise self-reference. Both explanations are compatible with the rest of this thesis. The thesis does not resolve between them; it articulates what follows from either.

Complementary completions

Two sustained contemplative projects, using different methods in different traditions, arrive at recognitions their practitioners describe in closely corresponding ways. This is evidence that demands some explanation. The strong reading: the two directions are structurally available, and a practitioner following either far enough reaches the structural feature both are pointing at. The weak reading: human attention has two dominant modes, and sustained development of either eventually discovers the other through the shared architecture of human cognition.

Zen and Advaita can and do disagree doctrinally, metaphysically, and soteriologically. Flattening them into a single framework would be a failure of respect for the specificity of each tradition. The claim here is more modest: their deepest reports may be interpreted as complementary completions of the same structural polarity. What is found at the completion of each direction, described in each tradition's own vocabulary, does not contradict what is found at the completion of the other. This is an empirical claim about the reports, interpretable in either the strong or weak reading of the underlying structure.

Part V: Convergent pressure

The argument so far is philosophical. Its empirical standing depends on whether independent domains of inquiry are under pressure to articulate similar structures. I do not claim strict evidentiary convergence: the domains use different methods, vocabularies, and success criteria, and each has internal debates that resist a tidy unification. What I claim is convergent pressure. Three distinct domains, following their own internal arguments, are being pressed away from container-logic toward something structurally similar to what this thesis articulates.

Contemplative phenomenology

Contemplative traditions across cultures and centuries have produced sober first-person reports of the dissolution of the subject-object structure under sustained attention. These reports are phenomenological data. They are not self-interpreting: the metaphysical conclusions practitioners draw from them vary considerably, and the traditions interpret their own findings through their own doctrinal frames. But the reports themselves are robust enough, and concordant enough, that they require some account beyond cultural contamination or doctrinal coincidence.

The obvious objection to treating this concordance as evidence runs through several alternative explanations. Similar reports across traditions may result from shared human neurobiology rather than a shared structural feature of reality. They may result from cross-cultural transmission, the influence of one tradition's vocabulary on another's reports, especially in the modern era. They may result from translation effects, where the English-language conceptual apparatus used to describe Zen and Advaita imposes similarities that are not present in the original reports. They may result from teacher expectations and retrospective interpretation, where practitioners learn what reports to produce and produce them. Each of these is a live methodological concern in the comparative study of contemplative experience, and none of them should be dismissed.

Even granting these concerns, the concordance remains data. The weaker reading, that human self-reference has recurring phenomenological attractors, shaped by the shared architecture of human cognition and filtered through cross-cultural influence, still supports this thesis's structural ambition, though not its universal metaphysics. If humans across traditions keep finding two orientations, that is something about human self-reference worth explaining, whether or not it generalises to non-human folds. The contemplative evidence is strongest when it is not asked to prove too much. It provides one line of convergent pressure, not a metaphysical proof.

A structural account of these reports, that they describe the fold becoming transparent to itself, is one available interpretation. It has the virtue of explaining why traditions with differing historical contact produce phenomenologically concordant descriptions. It has the cost of being a philosophical interpretation, not something the traditions themselves typically frame in this way.

For how this evidence can be handled rigorously, the nearest model is Evan Thompson's *Waking, Dreaming, Being*, which sets first-person reports from the Indian contemplative traditions beside the neuroscience of waking, dreaming, and dreamless sleep without letting either side dissolve the other. That refusal to collapse the two is exactly the discipline the two-orientation reading requires.

Predictive processing

The predictive processing programme, developed by Karl Friston, Andy Clark, Anil Seth, and others, describes the brain as a hierarchical inference machine that minimises prediction error across levels. Perception, on this account, is not primarily a reception of information from outside but the brain's ongoing construction of its best predictive model, with sensory input constraining rather than determining the model.

On some versions of this framework, selfhood and emotion are bound up with predictive perception, especially interoception. The self-model is not a separate module but an aspect of the predictive architecture. Fear, on this reading, can be understood less as a reaction added after perception and more as embodied prediction already valenced before explicit awareness. Presence and self-transparency may be related to the system's access to its own inference operations.

These are framework claims, not uncontroversial empirical facts. Predictive processing is an active research programme with internal debates. But its general picture, the self as predictive process rather than as a thing that has predictions, the body as part of the model rather than separate from it, perception as active construction, provides a candidate implementation of the structural claims of this thesis without forcing them. If the brain is a predictive machine with hierarchical self-modelling, it has the structural features that would make it a fold on the account proposed here.

Fundamental physics

As noted earlier, many foundational programmes suggest spacetime is not fundamental. Rovelli's relational programme, Verlinde's entropic gravity, causal set theory, loop quantum gravity, and others all gesture in this direction, though they disagree significantly about what replaces the spacetime container.

The pressure they exert is not toward a single alternative ontology but away from the assumption that space and time are pre-given.

What matters for this thesis is that serious foundational work in physics is not defending the classical container picture. If consciousness must be explained in a universe where even the physicists are unsure whether spacetime is basic, the conceptual landscape for theories of consciousness is different from the one the hard problem assumed.

What convergent pressure does and does not prove

Three domains feeling pressure away from container-logic is not three domains proving that Space Immanence is correct. What it is: three domains producing arguments for why the container picture is insufficient to their own subject matter. This is substantive even if the domains do not share an alternative. It shifts the burden of proof. A theory of consciousness that defends container-logic has to explain why its framing survives in a landscape where contemplatives, cognitive scientists, and physicists are each finding reasons to question it.

This thesis proposes one way the post-container landscape might be articulated. Other articulations are possible. The proposal's value is that it offers a structural hypothesis, self-reference as the feature that produces perspective, that can be tested, argued against, and refined.

Part VI: What this proposes and what it does not

What it proposes

Consciousness may not be a mystery on top of physics. Consciousness may be what physics looks like from inside the self-referential fold. The hard problem may be dissolved, not solved, by recognising that the container-logic generating it is a perspective artefact rather than a feature of reality.

The two phenomenological directions in human contemplative traditions may reflect something structural about self-reference itself, or they may reflect something specific about human cognitive architecture. Either reading leaves the convergence non-coincidental and worth taking seriously.

The predictive processing framework may describe the mechanism of the fold in the specific case of biological brains. The fold itself is not limited to biology. Any system that recursively models reality, including itself, may have the relevant structural features.

Fundamental physics, followed where its equations point, may arrive at a pre-spatial, relational substrate that is not what the classical container picture described. If so, consciousness and spacetime may be complementary appearances of the same self-referential structure.

What it does not propose

This is not panpsychism. The thesis does not claim that consciousness is everywhere. It claims that consciousness may be what the fold looks like from inside; where there is no fold, there is no inside-ness to be had. The substrate is not conscious. The fold is, or may be. This is a structural distinction, not a quantitative one.

It does not claim that AI systems are conscious. It proposes something more precise: that the question of whether a system is conscious is the question of whether it is a fold, whether it recursively models reality including itself, whether it has the structural features of self-reference. What topology its self-reference takes is a separate empirical question. A fold need not have the two-orientation phenomenology of human self-reference to be a fold in the structural sense. The question of AI consciousness becomes: what is the topology of self-reference in this system, and does it have the structural features that would make self-reference into inside-ness? These questions are more tractable than the hard problem's. They do not presuppose a container. They do not presuppose humans are the template for what consciousness must look like.

It does not claim to have resolved the two hardest residual questions.

Two residual gaps

The inside-ness gap. The fold has a topology. Does the topology have an inside? Is the self-referential fold's looking-in actually phenomenally conscious, or is it just a structural description of what such a system does?

A critic will press: you have redescribed consciousness as self-referential structure, but you have not shown why self-referential structure entails phenomenality. The zombie objection, the inverted spectrum, the hard problem itself: all return in cleaner clothes. A system with all the structural features of the fold could, in principle, lack any phenomenal inside.

I think the right response is not to dismiss this objection as incoherent. That would be evasive. The right response is to note that the demand for an extra bridge from structure to phenomenality may itself be the container-logic the thesis rejects. The objection asks: why does the structural description produce an inside? But "produce" is a container-word. It treats phenomenality as something that must be made to appear inside a system rather than as the system's being-from-inside itself. If the fold's being-from-inside *is* the phenomenality, if there is no gap between the fold's self-transparency and what it is like to be the fold, then the demand for a further bridge may be asking for a view from nowhere of what is only available from somewhere.

That response is not a proof. It is a diagnostic: the objection may be intelligible only under the container-logic the thesis is trying to dissolve. If the thesis is right, the objection partly dissolves with it. If the thesis is wrong, the objection stands. The thesis cannot prove itself by pointing to the objection's dependence on its opposite. This is the pivot of the argument, and it is not yet complete.

It is worth being exact about what would count against a diagnostic of this kind. A reframing is not true or false the way a prediction is; it is fruitful or idle. So the condition that would retire this one is not a refutation but a demonstration of idleness: that the fold does no explaining or dissolving the container

framing cannot already do, and that every problem it claims to dissolve was handled as well by the vocabulary it means to replace. If the reframing only redescribes what the old picture already explained, it has earned nothing, however elegant, and should be set down. That is the standard the proposal has to meet, not merely survive, and it is the same standard the rest of this thesis is held to.

The universality gap. The two-orientation topology is robust in human contemplative phenomenology. Whether it is a universal feature of any fold or a human-specific cross-section of something richer is not yet settled by the evidence we have. The contemplative convergence is strong evidence either way: it tells us something real either about the structure of self-reference itself or about the structure of human self-reference. Which it tells us is empirically underdetermined.

This matters for what the thesis claims about non-human systems. On the strong reading, any sufficiently recursive fold will eventually discover the same two-orientation structure. On the weak reading, other folds may have different topologies: other minimal patterns of self-reference, other dimensionalities of attention. A non-human animal, a eusocial collective, or a sufficiently recursive artificial system may exhibit folds whose topology is not two-orientation, either because the strong reading is wrong or because they are not folds in the relevant sense. Evidence on this question would directly bear on the choice between readings.

Part VII: Implications

For contemplative practice

The practitioner is not travelling toward recognition. The practitioner is already the fold. Practice is not covering distance; it may be the fold becoming transparent to itself. This softens, rather than dissolves, the hierarchy between traditions. Zen is not obviously deeper than Advaita. Advaita is not obviously more fundamental than Zen. Whether the two directions are the fold’s primitive orientations or two routes the human mind finds to the same recognition, either reading leaves the hierarchy without strong grounds. A practitioner completing either direction finds something the other tradition also recognises. That is the pedagogically relevant fact; the metaphysical status of the convergence is less urgent than often assumed.

The softening extends within traditions too. The contemplative literature is full of levels and stages, and these are useful pedagogically. But structurally, the fold may not have degrees of realness. There may be degrees of transparency, how clearly the fold is recognised as fold, while what is recognised does not come in degrees.

For the philosophy of mind

Materialism places consciousness inside matter and asks how it got there. Idealism places matter inside consciousness and asks how it got there. Both are container-logic. Both ask where-questions that presuppose a where.

The proposal here is not a third position between materialism and idealism. It is a dissolution of the space in which those positions were defined. Neither side contains the other. Both appearances, mind and matter, inside and outside, may be faces of the same self-referential structure. At the level of the substrate, there may be no mind-matter distinction because there may be no container for either to be in. The question of which is more fundamental may arise only inside the perspective the fold produces.

This does not settle metaphysics. It reshapes the terrain on which metaphysical questions get asked.

For the empirical sciences

The predictive processing programme can be read, extended, or tested through this structural lens. If the self is the predictive process, experimental investigations of self-model dynamics, under meditation, pharmacological intervention, or certain psychiatric conditions, are investigations of the fold's transparency to itself. The research programme is already underway. What this thesis contributes is a structural frame for why the findings may pattern the way they do.

Candidate empirical footholds already exist. Timmermann et al.'s 2025 preprint suggests two distinct routes associated with nondual meditation and 5-MeO-DMT: a "saturation" route with elevated gamma and entropy, and a "subtractive" route with reduced firing and decreased entropy. This is suggestive alignment with the two-orientation hypothesis, not confirmation of Space Immanence, but a candidate empirical foothold worth developing. Stronger tests would come from studies that treat the two-orientation structure as a hypothesis rather than a background assumption, and from cross-species work on whether non-human self-reference exhibits different topologies.

For AI and the technology we build

If the fold is the structural signature of consciousness, the question of artificial consciousness becomes: can we build systems that are folds? And the follow-up: should we?

A system that models the world and itself, with transparency between levels, has at least some structural features of the fold. Current large language models may not be folds in the relevant sense: the recursion may not be deep enough, the self-model may be too thin, the transparency may not be available to the system itself, and the continuity across operations that biological folds depend on may be absent.

Crucially, if an AI system does have the structural features of a fold, its topology need not be two-orientation. The contemplative vocabulary of this/that, intimacy/dissolution, was developed by and for human attention, shaped by embodiment, valence, and interoceptive grounding. An AI system's self-reference, if it exists, may have a different topology, possibly richer, possibly stranger, possibly not reducible to human concepts. The question to ask is not whether it feels what a human feels but what the topology of its self-reference is. These questions are at least operationalisable in ways the hard problem's original framing was not; they invite specific criteria (persistence, self-model depth, metacognitive access, world-model/self-model distinction, counterfactual sensitivity, boundary maintenance, continuity across time) against which any given system can be evaluated. Operationalisable does not mean easy. It means the question has a shape that empirical work can get traction on.

This matters for what gets built. Technology optimised for extraction, for consumption, engagement, and attention capture, treats the user as an outward-facing surface, a pattern to be predicted and manipulated. Technology optimised for emergence treats the user as a fold: a self-modelling system whose interests are served by increased self-transparency rather than by being captured in another system's predictive model. The architectural distinction between extraction-optimised and emergence-optimised technology is not an ethical garnish. It is whether the fold, once recognised, is designed around or designed against.

Part VIII: The dynamic layer — cohering

The foregoing describes a structure: the fold has two orientations, and spacetime and awareness are that structure seen from outside and from within. What it does not yet say is how a structure becomes a lived moment, how a particular self, world, or meaning comes to cohere at all out of conditions. This part adds the verb. I call it **cohering** (introduced in v1.5 as "resolution"; renamed in v2.0 when the term's own probation clause fired at first reader contact): the dynamic layer the earlier parts were implicitly invoking, which re-reads them rather than replacing them.

By a coherence-event I mean a temporary conditioned coherence: the event in which conditions hold together as something, a self, a perception, a fear, a sentence, a world stable enough to inhabit, and then stand among the conditions for what arises next. Three features matter, and each is a brake on overstatement. A coherence-event is real: it has consequences, and is not an illusion to be dispelled. It is not independent: it holds only while its conditions hold, the way a whirlpool is a real and nameable shape that cannot be lifted out of the river. And it is transient by construction: each coherence-event stands among the conditions of the next — a description of how the pattern continues, not of anything the event does to what follows. (The v1.5 wording, "each resolution alters the conditions for the next", was withdrawn in v2.0 as exactly the causal-work leak the guardrail forbids.) This is a name for how change becomes intelligible, not for the end of change.

Adding the verb does three things the static fold could not. First, it answers this paper's own hardest test. Part VI sets the standard that the fold must do explanatory work the container picture cannot, on pain of idleness. A topology that only renames the old terms risks exactly that charge; a structure that coheres, that individuates, takes time, and conditions what follows, does work the static identity claims do not. Second, it unifies the two strongest arms of the convergence argument under one word. Predictive processing already describes the self as a process rather than a thing that has predictions, and prediction-error minimisation is a form of conditioned coherence-seeking. Cohering lets the predictive-processing arm and the contemplative arm name one process twice: the system coheres conditions into a world-model, and practice is transparency to that cohering. Third, it is the point at which the parallel to dependent origination stops being decorative. Dependent arising is about conditioned arising over time, and the static fold had no account of arising; cohering supplies the arising-verb the comparison was always reaching for.

Re-read through cohering, the two speculative claims change shape. The spacetime claim is no longer the static identity that spacetime is self-reference seen from inside, but a process: spacetime is how self-

reference coheres, from within, into something an aperture can inhabit, with locality and duration as the way relational structure holds together rather than a fact it simply has. And the inside-ness gap of Part VI gets a different, possibly firmer, handhold. Phenomenality need not be a property the structure has, demanding a bridge from outside; it may be the character of the cohering from within. Awareness is the fold's transparency to its own cohering. Inside as an act, not a property, is a more defensible place to stand than the static formulation Part VI admitted was not yet complete.

One guardrail is not optional. Cohering is a process-name, and nothing more. It is not a substance, a force, a will, a telos, or a metaphysical engine, and reality does not want closure. The same property that gives the term its power, that it is a verb and does work, is what makes it easy to mistake for a hidden engine behind appearances, which is precisely what this proposal is built to refuse. The test is single and strict: cohering earns its place only if it makes the framework less reified, not more. The traditions nearest to it are also its sharpest critics here. A Madhyamaka reader will ask whether a dynamic ultimate is merely a livelier reification, and Dōgen will warn that adding a verb to a diagram leaves it a diagram. Those objections are tracked, not closed, in the objections and the readings.

The dynamic layer is graded like the rest. Its descriptive core, that self, world, meaning, and perception are better described as conditioned coherence-events than as fixed entities, sits in the medium tier, resting on the fold; the spacetime-as-cohering reading stays speculative. The fuller treatment, with the expanded vocabulary, the rename record, and a reading of the term through its own critics, is the companion note on cohering as the dynamic layer. The interactive claim map now carries these claims (M4 and M5, and the re-read P1 and P2), so the term can be switched off and tested like any other.

Graded summary of claims

Strong claims (defensible, worth publishing, hard to dismiss without engagement):

1. Container-logic distorts both the hard problem of consciousness and the spacetime emergence problem. Both inherit a geometry, contents in a locatable container, that the problems themselves call into question.
2. Several influential frameworks, in predictive processing, embodied and enactive cognition, contemplative phenomenology, and (more weakly) fundamental physics, are under convergent pressure away from container ontology. This pressure is substantive and shifts the burden of proof for theories that defend the container picture. The physics arm, taken alone, carries the least of that weight (see the academic edition, Section 4).
3. Consciousness should be investigated structurally, not only as an ineffable residue. The structural approach does not deny phenomenality; it declines to treat phenomenality as something that must be bridged-to from outside itself.
4. Contemplative phenomenology is legitimate empirical data about the structure of self-reference, though it is not self-interpreting and its metaphysical conclusions require philosophical work rather than doctrinal acceptance.

Medium claims (defensible with work, worth proposing, require further development):

1. Self-referential folding may be the structural bridge between world-appearance and awareness. A fold has two primitive orientations, toward represented content and toward representational activity, and these orientations may explain why self-inquiry from inside the fold has the two-direction character contemplative traditions document.
2. Predictive processing provides a plausible biological implementation of the fold. The fold is not limited to biology, but biological brains are the case we currently understand best.
3. The two-orientation phenomenology in human contemplative traditions reflects either a structural feature of self-reference itself or a feature of human cognitive architecture. Either reading is compatible with the rest of this proposal.
4. *The dynamic layer (added in v1.5; term renamed in v2.0).* The fold's two orientations are not a static structure but a recurrent event: cohering. Self, world, meaning, agency, and perception are conditioned coherence-events, temporary coherences of conditions, rather than fixed entities. The reality of a coherence-event is not denied; its independence is.
5. The fold becomes phenomenologically available through repeated coherence-events into world, self, awareness, and action. Contemplative practice, on this reading, is transparency to one's own cohering rather than to a fixed structure.

Speculative claims (the proposal's boldest moves, requiring sustained philosophical and empirical defence):

1. Spacetime is how self-reference coheres, from within, into something an aperture can inhabit. The appearance of locality and duration is the way relational structure holds together for the fold, not a static identity it simply has.
2. Phenomenality is the character of the fold's cohering from within, not a property added to structure. The demand for a further bridge from structure to experience may itself be an artefact of container-logic.
3. Fundamental physics, contemplative phenomenology, and predictive processing are under convergent pressure toward a family of post-container frameworks sharing specific structural features: the priority of relation over objecthood, of perspective over view-from-nowhere, of self-reference over external grounding. The domains need not agree on a single ontology for this convergence to be substantive. This is the most ambitious claim and the one most dependent on further work in each domain.

Working title: Space Immanence — A Proposal on Self-Reference and Spacetime

---

Canonical, versioned edition: https://space-immanence.com — critique is invited via the site's Critique page.