Distinction Before Objects

Projection, Reachability, and the Hidden Geometry of Reality

Flyxion

Independent Researcher

June 2026

Abstract. *Six research programmes developed over the past year—*RSVP *field theory,* CLIO *compression geometry, the Admissibility framework, Spherepop bubble calculus, the Repair architecture, and Distinguishability Geometry—appear at first to occupy different intellectual territories: physics, cognitive science, computation, formal ontology, memory theory, and information geometry respectively. This essay argues that they are not six theories but six coordinate systems on a single underlying object. That object is the* distinction structure of a representational situation*: the organised capacity to tell things apart, compress them, reach them, restore them, and express them. The convergence is not accidental. Each programme was driven independently toward the same inversion: from object-primary to distinction-primary ontology. What emerges is a unified geometry in which objects are not primitive but derived—residues left by the operation of distinction on raw possibility.*

Contents

1. The Convergence Problem

There is a certain moment in a research programme when you notice that the thing you have been building from one direction looks, from the right angle, exactly like the thing someone else has been building from another. The question is whether the resemblance is deep—a shared object seen from different vantage points—or superficial, a family of analogies that dissolves on close inspection.

This essay is concerned with a convergence of that kind, except that the "someone else" is the same researcher at different times, working on what appeared to be different problems. Over the past year, six frameworks were developed in sequence, each motivated by a distinct theoretical need:

- RSVP (Relativistic Scalar-Vector Plenum) was developed to give a field-theoretic account of physical possibility—to describe the plenum of states from which actual events are drawn.
- CLIO (Compression-Layer Information Ontology) was developed to model how cognition and language manage complexity through hierarchical projection—how a rich world becomes a tractable representation.
- The *Admissibility* framework was developed to ask which transitions in a representational or physical system are genuinely reachable—which moves remain within the constraint-compatible region.
- *Spherepop* was developed as a computation model in which the primitive operation is the collapse of a bubble—a history—into its result, with attention to what is preserved and what is destroyed at each pop event.
- The *Repair* architecture was developed to answer a question in the theory of memory and systems: when functionality is lost, what does restoration require, and what can never be recovered?
- *Distinguishability Geometry* was developed to give a formal account of the capacity to express distinctions—to make "what can be told apart" more primitive than "what exists."

Individually, each of these looks like a domain-specific contribution. RSVP is a

physics paper. CLIO is a cognitive science paper. Admissibility is a formal ontology paper. Spherepop is a programming languages paper. Repair is a systems paper. Distinguishability Geometry is an information theory paper.

This essay is an argument that they are the same paper, written six times, from six starting points, converging on a single inversion that none of them was initially designed to articulate.

That inversion is: *distinction before objects*.

2. The Standard Picture and Its Failure

The default ontology of almost every formal discipline begins with objects. In set theory, elements precede relations. In physics, particles or fields are primary; interactions are secondary. In logic, individual constants come first; predicates describe them. In cognitive science, concepts are taken to be representations of pre-existing categories; the categories are assumed to carve nature at its joints. In computer science, data structures hold objects; algorithms operate on them.

This picture has enormous practical utility. It is clean, composable, and aligns with the intuition that the world contains definite things that we then perceive, name, and reason about.

But it faces a systematic difficulty when applied to the study of *representation, compression, and change*. The difficulty is this: objects, as classically conceived, do not degrade gracefully. They are either present or absent. Information about them is either known or unknown. A representation either contains them or does not.

Real representational systems do not behave this way. They *confuse* objects—not randomly, but systematically, according to structural principles that the object-primary framework has no vocabulary to describe. A compression scheme does not lose an object; it *collapses a distinction*. A failing memory does not forget a fact; it loses the *differential profile* that made the fact distinguishable from neighbouring facts. A phase transition in a physical system does not destroy particles; it *equates states* that were previously different. A paradigm shift in science does not discard observations; it *reconstitutes the identity conditions* for what counts as the same

or different observation.

In each case, the vocabulary of objects is too coarse to describe what is actually happening. The natural vocabulary is the vocabulary of *distinctions*: what can be told from what, at what cost, and under what conditions.

Each of the six frameworks arrived at this conclusion independently, from the direction of its own motivating problems. The rest of this essay traces those six paths to their common destination.

3. Six Paths to the Same Inversion

3.1. rsvp*: The Field-Theoretic Path*

RSVP began as an attempt to give a substrate-independent account of physical possibility. The question it asked was: what is the structure of the space from which events are drawn? Not: what are the fundamental particles? Not: what are the laws? But: what is the *plenum*—the field of possibility that actual events are selections from?

The scalar component of the RSVP field encodes local state density: how many distinguishable states are available at a given point and scale. The vector component encodes directed flow: how possibility moves, propagates, and concentrates. A phase transition is a sharp change in the topology of the scalar field—a reorganisation of which states are available rather than which states are occupied.

What RSVP discovered, in the process of developing this account, is that the fundamental quantity is not the state itself but the *degree to which states are mutually distinguishable at a given scale*. The scalar field is not a density of objects; it is a density of *available distinctions*. Particles, fields, and events are high-distinction-density regions in the plenum. Phase transitions are distinction-density collapse events.

The path from RSVP to distinction-primary ontology runs through the question: what does it mean for two physical states to be "the same"? The answer, in RSVP terms, is not that they have the same properties but that they are *not distinguished by any observable at the current scale*. Identity is indistinguishability. Objects are

equivalence classes of indistinguishable states. The field comes first; the objects are what the field's distinction structure leaves behind.

3.2. clio*: The Representational Path*

CLIO began from the opposite end: not from physical possibility but from the problem of cognitive and linguistic compression. The question was: how do systems that cannot represent everything manage to represent anything usefully?

The answer CLIO gave was hierarchical projection. At each level of the hierarchy, a richer space is projected onto a coarser one. Information is lost in the projection, but the loss is structured: the system preserves what matters for the tasks downstream and discards what does not. The fibers of the projection—the sets of pre-images that collapse to a single image—are the basic units of compression.

As CLIO developed, it became clear that the central question was not "how much was compressed?" but "what distinctions were destroyed?" Compression is not primarily a volume-reduction operation; it is a distinction-collapsing operation. A projection is benign if it collapses states that were already equivalent for downstream purposes—states that the system had no use for distinguishing. It is harmful if it collapses states that differ in ways that matter later.

This led directly to the distinction-primary perspective: the content of a representation is not the objects it contains but the distinctions it preserves. Two representations are equivalent not when they describe the same objects but when they preserve the same distinctions. Compression quality is measured not by retained content but by retained distinction capacity.

CLIO names the projection; distinguishability geometry measures the damage. But the conceptual commitment is the same: distinctions are the fundamental currency.

3.3. Admissibility: The Reachability Path

The Admissibility framework developed from a question in formal ontology and process theory: given a current state and a proposed transition, is the transition *reachable*? Not merely possible in a logical sense, but reachable through a sequence of constraint-compatible steps from here?

The framework builds a geometry of reachability: the set of states accessible from a given starting point under a given constraint field. A transition is admissible if it remains within this reachable region; inadmissible if it requires a jump across a constraint boundary that cannot be traversed continuously.

What the Admissibility framework discovered in the course of this development is that the constraint field is itself a distinction structure. The boundaries that define the reachable region are precisely the points at which states become indistinguishable under the constraint: crossing the boundary is equivalent to losing a distinction that was previously maintained. Admissibility is, at bottom, a theory of *distinction preservation under transition*.

A trajectory is admissible if the distinctions it relies on are preserved throughout. A trajectory is inadmissible if it collapses distinctions in ways that the downstream process cannot recover from. The threshold conditions of the Admissibility subcategory—the ϵ-bounds on deficit increase—are the formal expression of this commitment: we permit transitions that do not destroy too much distinction capacity.

The path from reachability to distinctions is thus direct: you cannot navigate to a state you cannot distinguish from your current location. Reachability presupposes distinguishability. The admissibility geometry is a geometry of preserved distinctions.

3.4. Spherepop: The Computational Path

Spherepop was built as a computation model centred on the *bubble*: a structured region with an interior, a boundary, and a context. Computation is the organised collapse of bubbles—pop events that expose interior structure to the context—and the question the model was designed to answer was: what is preserved and what is destroyed by a pop?

The pop event looked, at first, like a purely operational notion: a bubble collapses when its boundary conditions are satisfied, and the computation continues. But the question of what is *preserved* under a pop turned out to be a question about distinctions. A bubble's interior may contain structure that is distinct from the context's current representation of it. A pop that exposes this structure reduces the

distinction deficit: the context can now tell apart things it previously could not. A collapse that fails—a bubble that cannot be popped because the boundary resists—is a failure of distinction expression: the language does not have the resources to make the interior distinctions available.

The grammar of Spherepop—the rules governing which bubbles can form, interact, and collapse—turned out to be exactly a description language $\mathcal{L}$ in the sense of distinguishability geometry. Pop is distinction revelation. Collapse is projection. Refusal is distinction preservation. Analogy between systems is transport with controlled distortion.

Spherepop thus arrived at the distinction-primary picture from the computational direction: the fundamental operation of computation is not the manipulation of objects but the management of distinctions—their revelation, compression, preservation, and transport.

3.5. Repair: The Restoration Path

The Repair framework addressed a question that is, on its face, the most practical of the six: when a system has lost functionality, what does genuine restoration require? Not the superficial restoration of outputs, but the deep recovery of the capacity to produce those outputs.

The answer Repair gave was that genuine restoration requires *re-distinguishing* what was conflated in the failure. A system that has lost functionality has, in almost every interesting case, lost the ability to distinguish states that its correct operation depends on separating. A damaged memory does not lack records; it lacks the differential profiles that made records identifiable. A corrupted inference engine does not lack logical steps; it has collapsed distinctions that the logical steps depend on.

This is why Repair turns out to require more than restoration of outputs: it requires restoration of *distinction structure*. A system whose output is repaired but whose internal distinctions remain collapsed will fail again, because the failure mode was not in the output but in the distinction architecture that generated it.

Repair names the operation that distinguishability geometry calls embedding with

revision: expanding the description language to recover distinctions that were projected away. The Repair framework thus arrives at the distinction-primary picture from the direction of failure: what breaks first is not objects but the differences between them.

3.6. Distinguishability Geometry: The Invariant Path

Distinguishability geometry was developed last, partly in response to noticing the convergence described above. It begins with the explicit decision to make distinction capacity primitive: the ontological triple $(X, \sim, \mathcal{L})$ encodes a space of elements, an indistinguishability relation, and a description language that bears the cost of expressing distinctions.

What distinguishability geometry adds that the other five frameworks do not is the *invariant*: the pair of deficit measures $(\delta_{\text{code}}, \delta_{\text{dist}})$ that quantifies, for any representational situation, how far it falls from the ideal of perfect distinction expression. The four operations (projection, embedding, revision, transport) are the dynamics; the deficit is what they move.

The contribution of distinguishability geometry to the unified picture is precisely this: it supplies the *common language* in which the other five frameworks can communicate. RSVP field evolution changes the deficit; CLIO projections increase it; admissibility bounds it; Spherepop's pops and collapses move it locally; Repair's restorations decrease it. The deficit is the invariant that travels across all five contexts unchanged in meaning, even as the physical, computational, and cognitive details vary.

4. The Common Object

The claim of this essay is that the six frameworks are coordinate systems on a single object. What is that object?

It is not a theory, exactly—at least not in the sense of a fixed set of axioms with a fixed domain of interpretation. It is more like a *geometric structure*: a space equipped with an organised capacity to distinguish its points, a dynamics that moves through that space by collapsing and recovering distinctions, and a measure that tracks how much distinction capacity is available at each point.

Call it the *distinction field* of a representational situation. A representational situation is anything that has states, a way of confusing or separating those states, and a language for expressing whatever separations are available. Physical systems are representational situations. Minds are representational situations. Computations are representational situations. Memories, social institutions, and scientific theories are representational situations.

The distinction field of such a situation has:

- A *topology* given by the indistinguishability relation: which states are confusable and which are not, at what scales, under what operations.
- A *dynamics* given by the four operations (projection, embedding, revision, transport) that move the situation through its possible configurations.
- A *measure* given by the deficit: how far the current description language falls from expressing all available distinctions.
- A *constraint* given by the admissibility condition: which transitions are reachable without catastrophic distinction loss.

RSVP describes this structure in the language of physical fields. CLIO describes it in the language of projection hierarchies. Admissibility describes it in the language of reachability constraints. Spherepop describes it in the language of computational bubble dynamics. Repair describes it in the language of recovery from collapse. Distinguishability geometry describes it in the language of information theory and category theory.

None of these descriptions is more fundamental than the others. Each is natural for its domain. Each illuminates aspects that the others make less visible. But the object they describe is the same.

4.1. Objects as Residues

The title of this essay—"Distinction Before Objects"—makes a substantive claim, not merely a methodological one. It is not only saying that it is *useful* to think of distinctions before objects. It is saying that objects are *derived* from distinctions, not the other way around.

The derivation goes like this. Begin with a space of states and an indistinguishability relation. The equivalence classes of the relation are the *observable states*: clusters of indistinguishable elements that function, from the outside, as units. These clusters are what we call objects.

An object, on this account, is not a primitive. It is a *residue of distinction*: the trace left by the operation of an indistinguishability relation on a richer underlying space. When the relation is coarsened, objects merge. When it is refined, objects split. The objects we experience are the objects given by the indistinguishability relation that our current representational apparatus imposes on the world.

This is not idealism. The underlying space of states is real; it is not constituted by the relation. But the *individuation* of that space into objects—the carving of it into identifiable, nameable, persistent things—is a relational achievement, not a primitive fact. Objects are not given; they are constructed by the distinction structure of the situation we bring to bear on the world.

The consequence is that the question "how many objects are there?" is always relative to a distinction structure. It is not a question with an absolute answer. What has an absolute answer is: what is the distinction capacity of the current situation, and how far is it from being fully expressed?

5. The Inversion and Its Consequences

The object-primary picture and the distinction-primary picture are not merely different descriptions of the same ontology. They lead to different research programmes, different notions of failure, different standards for what counts as a good explanation, and different pictures of what progress looks like.

5.1. Different Notions of Failure

In the object-primary picture, a system fails by losing objects—by forgetting facts, corrupting records, misidentifying individuals. The repair strategy is restoration of the lost objects: recovery of the forgotten facts, correction of the corrupted records.

In the distinction-primary picture, a system fails by losing *distinctions*—by col-

lapsing states that should have been kept separate, by allowing the description language to fall out of alignment with the distinction structure of the domain. The repair strategy is not restoration of objects but *recovery of distinction capacity*: re-expansion of the language, re-refinement of the relation, re-embedding of the system in a richer space where the collapsed states can be separated again.

This is not a small difference. The object-primary failure mode is *loss*; the distinction-primary failure mode is *collapse*. Loss is addressed by recovery; collapse is addressed by re-distinction. A system that has lost an object can in principle be restored by finding that object again. A system that has collapsed a distinction cannot be restored by finding an object, because the problem is not that an object is missing; it is that the *boundary* between two objects has been erased.

5.2. Different Standards for Explanation

In the object-primary picture, a good explanation identifies the objects responsible for an effect: the particle, the gene, the agent, the mechanism. Explanation is complete when the relevant objects and their properties have been specified.

In the distinction-primary picture, a good explanation identifies the *distinction structure* responsible for an effect: which distinctions were available, which were expressed, which were collapsed, and what the deficit was. Explanation is complete when the distinction dynamics have been traced.

This is a richer standard. It asks not only "what was there?" but "what could be told from what?"—and that question, as RSVP, CLIO, Admissibility, Spherepop, Repair, and Distinguishability Geometry all discovered, is often the more fundamental one.

5.3. Different Pictures of Progress

In the object-primary picture, theoretical progress is accumulation: more objects discovered, more properties specified, more mechanisms identified. The world grows by the addition of new entities.

In the distinction-primary picture, theoretical progress is *refinement*: the distinction structure of the current theory is improved, its deficit is reduced, its language is brought into closer alignment with the distinctions the domain actually makes.

The world does not grow; our *resolution* of it improves. Progress is not addition but clarification.

This picture is more consonant with the history of science as actually practised. Major theoretical advances—the Copernican revolution, the Darwinian synthesis, the quantum mechanical reconstitution of matter—are not primarily additions of new objects. They are reconstitutions of *identity conditions*: revisions of which things count as the same or different, which distinctions matter and which do not. Copernicus did not add planets; he changed the distinction structure governing planetary and stellar motion. Darwin did not add species; he redrew the distinction between species and variety. Quantum mechanics did not add particles; it made some classical distinctions (position, momentum simultaneously) inadmissible.

6. The Unified Geometry

We are now in a position to state the unified geometry more precisely.

Central Claim. The six frameworks—RSVP, CLIO, Admissibility, Spherepop, Repair, and Distinguishability Geometry—are coordinate systems on the *distinction field* of a representational situation. The distinction field consists of a topology (the indistinguishability relation), a dynamics (the four operations), a measure (the deficit), and a constraint (the admissibility condition). Objects are residues of this field, not its primitives.

The claim can be tested. If it is correct, then every concept from each framework should have a natural translation into every other framework. There should be no concepts that are irreducibly local to one framework. Let us verify this for a selection of core concepts.

6.1. Cross-Framework Dictionary

Projection. In CLIO, projection is the operation that compresses a representation by mapping a richer space to a coarser one. In RSVP, projection corresponds to scale-reduction in the scalar field: the integrated-out degrees of freedom become unavailable. In Admissibility, a projection that crosses a constraint boundary is inadmissible; one that stays within the constraint field is admissible. In Spherepop, collapse is projection: a bubble's interior is merged with its context, and the

boundary is destroyed. In Repair, projection is the failure mode: it is what needs to be undone when functionality is lost. In Distinguishability Geometry, projection is the first of the four operations: it coarsens the indistinguishability relation and typically increases the distinguishability deficit.

Recovery/embedding. In CLIO, embedding is the inverse of compression: moving from a coarser to a richer representation, recovering distinctions. In RSVP, embedding is expansion into a richer phase space: new degrees of freedom become available. In Admissibility, an embedding expands the reachable region. In Spherepop, push events (formation of new bubbles) are embeddings: new structure is made available. In Repair, embedding is the core operation: re-expansion into a richer language that can separate states previously conflated. In Distinguishability Geometry, embedding is the second operation: it refines the relation and decreases the deficit.

Phase transition / paradigm shift. In RSVP, a phase transition is a sharp reorganisation of the scalar field's topology. In CLIO, a paradigm shift is a revision of the projection hierarchy itself—not a change in what is projected, but in how the levels are defined. In Admissibility, a phase transition may move the system to a new constraint region with different admissibility conditions. In Spherepop, a grammar revision changes which bubbles can form—a structural shift in the computation space. In Repair, a phase transition may make previously recoverable distinctions permanently inaccessible, or may open up new repair possibilities. In Distinguishability Geometry, a phase transition is a revision: $(X, \sim, \mathcal{L}) \mapsto (X, \sim, \mathcal{L}')$ with a possibly discontinuous change in deficit.

Information loss. In CLIO, information loss is projection distortion: the size of the fibers that are collapsed. In RSVP, information loss is a decrease in the distinction density of the scalar field. In Admissibility, information loss that exceeds the admissibility threshold is inadmissible—the system has left the reachable region. In Spherepop, information loss is a collapse without a corresponding push: structure is destroyed without being made available to the context. In Repair, information loss is precisely what repair addresses: the distinctions that were lost

are what the repair operation must recover. In Distinguishability Geometry, information loss is the increase in the distinguishability deficit: the gap between available and expressed distinctions widens.

6.2. The Invariant

The cross-framework dictionary confirms that there is a single invariant travelling through all six accounts: the distinction structure of the current situation, measured by the deficit and governed by the admissibility constraint.

This invariant can be stated simply: *how many distinctions are available, how many are expressed, and how far can we go without losing more than we can afford?*

RSVP answers this question in the language of field density and topology. CLIO answers it in the language of projection hierarchies and fiber entropy. Admissibility answers it in the language of reachable regions and constraint compatibility. Spherepop answers it in the language of bubble dynamics and pop events. Repair answers it in the language of restoration capacity and irreversible collapse. Distinguishability Geometry answers it in the language of the ontological triple and the deficit measures $(\delta_{\text{code}}, \delta_{\text{dist}})$.

The answers are the same answer.

7. Where This Sits in the Larger Programme

This essay is intended to occupy a specific position in the body of work it synthesises. *Constraint Before Content* established the foundational commitment to process ontology and constraint-first thinking. *The Admissibility Field* gave the most technically developed formulation of admissibility geometry. The present essay sits between them: more philosophical than the former in its willingness to draw out the ontological consequences of the convergence, but less formalism-heavy than the latter in its primary orientation toward the conceptual unification.

Its role in the programme is to give Distinguishability Geometry a larger function than that of a standalone framework. Distinguishability Geometry is not the seventh framework in a series; it is the framework that explains why the other six are related. It supplies the common language—the invariant, the measure, the

categorical structure—that makes the convergence legible.

The claim is not that Distinguishability Geometry is more fundamental than the others. That would be to repeat the error of object-primary ontology at the framework level: elevating one perspective to a privileged position and treating the others as its derivatives. The point is precisely that none of the six is more fundamental; they are six good coordinate systems for a single object, and the right choice of coordinates depends on the problem at hand.

What Distinguishability Geometry provides is not primacy but *translatability*: the common vocabulary in which results from any one framework can be understood by any of the others.

7.1. Open Commitments

The unified geometry carries several open commitments that future work must address:

The Classification Conjecture. Every representational transformation factors into a finite composition of projections, embeddings, revisions, and transports. This is the claim that the four operations are complete. Its proof would close the geometry; its failure would identify a fifth operation and a gap in the unified picture.

Deficit Equivalence. Zero coding deficit and zero distinguishability deficit coincide for well-behaved triples. If true, a single scalar suffices as the representational invariant. If false, the two deficits are genuinely independent quantities, and the geometry is richer than currently described.

The rsvp Derivation. The deficit landscape should be derivable analytically from RSVP field equations. This would close the loop between the physical and informational descriptions, making the correspondence between field dynamics and distinction geometry explicit and computable.

Completeness of Repair. The Repair architecture maintains a ledger of collapse events. The question is whether this ledger is complete in the distinction-

geometric sense: does it record enough to reconstruct the full distinction trajectory of the system? Completeness here would mean that nothing about the system's distinction history is inaccessible from its repair record.

Spherepop Soundness. Is there a Spherepop operational semantics that is sound with respect to distinguishability geometry—every reduction step either preserves or improves the deficit? Such a semantics would make Spherepop a *certified* distinction-preserving computation model.

8. Conclusion: The World Before Objects

The title promised a claim: distinction before objects. The essay has tried to earn it.

The claim is not that objects do not exist. Tables, particles, persons, and propositions exist well enough for most purposes. The claim is that their existence is *derived*: they are the residues left by the operation of a distinction structure on a richer underlying space. Change the distinction structure—refine it, coarsen it, revise it, transport it to a new domain—and the objects change with it. Not because the world has changed, but because the boundary conditions under which states count as the same or different have changed.

This is the inversion that RSVP, CLIO, Admissibility, Spherepop, Repair, and Distinguishability Geometry all arrived at, from six different directions, over the course of a year of independent development. The convergence is the evidence. Six coordinate systems on one object are more reliable witnesses than one.

The geometry that unifies them is not a finished theory. The four open commitments named above are genuine open problems, not rhetorical gestures. The cross-framework dictionary, while compelling, requires formal verification in each case. The claim that objects are residues of distinctions, while supported by the six-fold convergence, has not been established beyond philosophical argument.

But the direction is clear. The world before objects is a world of distinction fields, deficit landscapes, and reachable transformations. Objects emerge from it the way peaks emerge from a topography: not as primitives, but as the places where

the distinction density is highest, the deficit is lowest, and the admissible paths converge.

* * *

The strongest summary this programme can currently offer is this: RSVP describes the field dynamics of possibility; CLIO describes projection into manageable representations; Admissibility decides which transformations are allowed; Spherepop operationalises local distinction events; Repair records their collapse residues and restores what can be restored; and Distinguishability Geometry supplies the invariant that lets all five talk to each other. Together they describe a single thing: the organised capacity of a situation to distinguish its own states, and the damage it sustains, and the recovery it achieves, in the course of being a situation at all.

Flyxion is an independent researcher working at the intersection of theoretical physics, philosophy of computation, and cognitive science. This essay synthesises work across the RSVP field theory, CLIO compression geometry, Admissibility framework, Spherepop computation model, Repair architecture, and Distinguishability Geometry programmes developed over 2025–2026.